EFI 1.10
Driver Writer's Guide

Draft for Review

Version 0.9

July 20, 2004

Acknowledgements

This document was developed in close consultation with Hewlett-Packard. Hewlett-Packard has made substantial contributions to the content presented here based on their experience with EFI 1.10 driver development, integration, validation, and testing. These efforts are gratefully acknowledged.

Revision History

Revision	Revision History	Date
0.31	Initial draft.	4/3/03
0.70	Initial draft. Edited for formatting and grammar.	6/3/03
0.9	Incorporated industry review comments. Added new chapters for USB and SCSI. Updated the coding conventions in chapter 3. Updated for the 1.10.14.62 release of the EFI Sample Implementation. Updated the supported versions of Microsoft Visual Studio* and Windows*. Removed TBD chapters that appeared in the 0.7 version. Edited for grammar and formatting.	7/20/04

Contents

3 Coding Conventions

4 EFI Services

5 General Driver Design Guidelines

6 Classes of EFI Drivers

7 Driver Entry Point

8 Private Context Data Structures

16 SCSI Driver Design Guidelines

17 Driver Optimization Techniques

18 Itanium Architecture Porting Considerations

19 EFI Byte Code Porting Considerations

20 Building EFI Drivers

21 Testing and Debugging EFI Drivers

Figures

Tables

Examples

1 Introduction

1.1 Overview

This document is designed to aid the development of EFI drivers that follow the EFI Driver Model that is described in the *Extensible Firmware Interface Specification,* version 1.10 (hereafter referred to as the *"EFI 1.10 Specification"*). There are several different classes of EFI drivers and many variations of each of them. This document provides basic information for some of the most common classes of EFI drivers. Many other driver designs are possible. In addition, the design guidelines for the different driver-related protocols are covered, along with the design guidelines for PCI, USB, and SCSI buses. Finally, porting considerations for Itanium®-based platforms and EFI Byte Code (EBC) drivers and driver optimizations techniques are discussed.

This document assumes that the reader is familiar with the following:

- *EFI 1.02 Specification*
- *EFI 1.10 Specification*
- *EFI 1.10.14.62 Sample Implementation*

The *EFI 1.10.14.62 Sample Implementation* is also referred to as the *EFI Sample Implementation* or the *EFI 1.10 Sample Implementation* throughout the remainder of this document. The *EFI Sample Implementation* supports the following operating systems:

- Microsoft Windows NT* 4.0
- Microsoft Windows* 2000
- Microsoft Windows XP

It has a build infrastructure that supports Intel® compilers and the following versions of Microsoft Visual Studio* or Visual C++*:

- Microsoft Visual Studio .NET 2003 or 2002. These versions do not require service packs. It is recommended to use the Visual Studio .NET versions instead of Visual C++ 6.0.
- Microsoft Visual C++ 6.0 with the Visual Studio 6.0 Service Pack 5 or later and the Visual C++ Processor Pack Download. These two service packs are available for download from:

  http://msdn.microsoft.com/vstudio/downloads/updates/sp/vs6/sp5/default.aspx

  http://msdn.microsoft.com/vstudio/downloads/tools/ppack/download.aspx

In general, this document will use the term *Visual Studio* to refer generically to any of the supported versions in the Visual Studio or Visual C++ tool chains. If required, the *EFI Sample Implementation* can be modified to support a wider variety of operating systems and tools chains. However, only the supported operating systems and tools chains will be discussed in this document as required.

1.2 Organization of This Document

This document is not intended to be read front to back. Instead, it is designed more as a cookbook for developing and implementing drivers. The first four chapters provide background information, chapter 5 provides the basic recipe for all drivers, and the remaining chapters provide detailed information for developing specific types of drivers.

In general, driver writers should use this document in the following way:

1. Read the first four chapters of this document before starting to code.
2. Read chapter 5, which describes the general guidelines for designing all types of EFI drivers. Specifically, see section 5.2 for the general steps to follow when designing a driver. This section then points you to other sections in this document that contain the specific "recipe" for that particular type of EFI driver.
3. Read the specific sections or chapters listed in section 5.2 that apply to your EFI driver.

Table 1-1 describes the organization of this document.

Table 1-1. Organization of the EFI 1.10 Driver Writer's Guide

Chapter	Description
1. Introduction	Provides a brief overview of this document, its organization, and the conventions used in it.
2. Foundation	Describes the basic concepts in the *EFI 1.10 Specification.*
3. Coding Conventions	Describes the common coding conventions to use in an implementation of the *EFI 1.10 Specification.*
4. EFI Services	Describes the services available in the *EFI 1.10 Specification* and provides example code that uses these EFI services.
5. General Driver Design Guidelines	Gives general guidelines for designing all types of EFI drivers. Provides the basic "recipe" for developing a driver and points to later chapters or sections for more detailed information.
6. Classes of EFI Drivers	Describes the required and optional features and the subtypes for different types of drivers that follow the EFI Driver Model.
7. Driver Entry Point	Describes the entry point and optional features for the different classes of EFI drivers.
8. Private Context Data Structures	Introduces object-oriented programming techniques for managing controllers and the concept of private context data structures.
9. Driver Binding Protocol	Describes the requirements and features for the Driver Binding Protocol. This protocol provides services that can be used to connect a driver to a controller and disconnect a driver from a controller.
10. Component Name Protocol	Describes the requirements and features for the Component Name Protocol. This protocol provides a human-readable name for drivers and the devices that drivers manage.

continued

Table 1-1. Organization of the EFI 1.10 Driver Writer's Guide (continued)

Chapter	Description
11. Driver Configuration Protocol	Describes the requirements and features for the Driver Configuration Protocol. This protocol allows the user to set the configuration options for the devices that drivers manage.
12. Driver Diagnostics Protocol	Describes the requirements and features for the Driver Diagnostics Protocol. This protocol allows diagnostics to be executed on the devices that drivers manage.
13. Bus Specific Driver Override Protocol	Describes the requirements and features for the Bus Specific Driver Override Protocol. This protocol matches one or more drivers to a controller. In general, this protocol applies only to bus types that provide containers for drivers on their child devices.
14. PCI Driver Design Guidelines	Describes the guidelines that apply specifically to the management of PCI controllers.
15. USB Driver Design Guidelines	Describes the guidelines that apply specifically to the management of USB controllers.
16. SCSI Driver Design Guidelines	Describes the guidelines that apply specifically to the management of SCSI controllers.
17. Driver Optimization Techniques	Describes several techniques to optimize an EFI driver.
18. Itanium® Architecture Porting Considerations	Describes guidelines to improve the portability of an EFI driver and the pitfalls that may be encountered when an EFI driver is ported to an Intel® Itanium® processor.
19. EFI Byte Code Porting Considerations	Describes considerations when writing drivers that may be ported to EFI Byte Code (EBC).
20. Building EFI Drivers	Describes how to write, compile, and package EFI drivers for the *EFI 1.10 Sample Implementation* environment.
21. Testing and Debugging EFI Drivers	Describes techniques to test and debug EFI drivers.
A. EFI Data Types	Lists the set of base data types that are used in all EFI applications and EFI drivers and the modifiers that can be used in conjunction with the EFI data types.
B. EFI Status Codes	Lists the **`EFI_STATUS`** code values that may be returned by EFI Boot Services, EFI Runtime Services, and EFI protocol services.
C. Quick Reference Guide	Provides a summary of the services, protocols, macros, constants, and GUIDs that are available to EFI drivers.
D. Disk I/O Protocol and Disk I/O Driver	Provides the source files to the Disk I/O Protocol, the EFI Global Variable GUID, and the source files to the disk I/O driver.
E. EFI 1.10.14.62 Sample Drivers	Lists all the sample drivers that are available in the *EFI Sample Implementation.*
F. Glossary	Lists and defines the terms and acronyms used in this document.

1.3 Related Information

The following publications and sources of information may be useful to you or are referred to by this specification:

- *Extensible Firmware Specification,* ver. 1.02, Intel Corporation, 2001, http://developer.intel.com/technology/efi
- *Extensible Firmware Specification,* ver. 1.10, Intel Corporation, 2002, http://developer.intel.com/technology/efi
- *EFI Driver Library Specification,* ver. 1.11, Intel Corporation, 2003, http://developer.intel.com/technology/efi
- *EFI 1.10.14.62 Sample Implementation*, Intel Corporation, 2003, http://developer.intel.com/technology/efi
- *EFI 1.1 Shell Commands Specification*, ver. 0.3, Intel Corporation, 2003, http://developer.intel.com/technology/efi
- *Itanium® Processor Family System Abstraction Layer Specification,* Intel Corporation, http://developer.intel.com/design/itanium/family
- *Intel® Itanium® Architecture Software Developer's Manual,* vols. 1–4, Intel Corporation, http://developer.intel.com/design/itanium/family
- *A Formal Specification of Intel® Itanium® Processor Family Memory Ordering,* Intel Corporation, http://developer.intel.com/design/itanium/family
- *Microsoft Portable Executable and Common Object File Format Specification*, Microsoft Corporation, http://www.microsoft.com/whdc/system/platform/firmware/PECOFF.mspx
- *Developer's Interface Guide for 64-bit Intel Architecture-based Servers,* ver. 2.1, Compaq Computer Corporation, Dell Computer Corporation, Fujitsu Siemens Computers, Hewlett-Packard Company, Intel Corporation, International Business Machines Corporation, and NEC Corporation, 2001, http://www.dig64.org/specifications
- *Code Complete*, Steven C. McConnell, ISBN 1-55615-484-4

1.4 Conventions Used in This Document

This document uses the typographic and illustrative conventions described below.

1.4.1 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 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.

1.4.2 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*.

1.4.3 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 electronic version of this specification, any plain text underlined and in blue indicates an active link to the cross-reference.
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.
`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 and can be outlined by a thin black border.

2
Foundation

There are several EFI concepts that are cornerstones for understanding EFI drivers. These concepts are defined in the *EFI 1.10 Specification*. However, programmers who are new to EFI may find the following introduction to a few of the key EFI concepts a helpful framework to keep in mind as they study the *EFI 1.10 Specification*.

The basic concepts that are covered in the following sections include the following:

- Objects managed by EFI-based firmware
- EFI System Table
- Handle database and protocols
- EFI images
- Events
- Device paths
- EFI Driver Model
- Platform initialization
- Boot manager and console management

As each concept is discussed, the related application programming interfaces (APIs) will be identified along with references to the related sections in the *EFI 1.10 Specification*. One component that is distributed with the *EFI 1.10 Sample Implementation* is the EFI Shell. The EFI Shell is a special EFI application that provides the user with a command line interface. This command line interface provides commands that are useful in the development and testing of EFI drivers and EFI applications. In addition, the EFI Shell provides commands that can help illustrate many of the basic concepts that are described in the following sections. The useful EFI Shell commands will be identified as each concept is introduced. The EFI Shell commands are described in detail in the *EFI 1.1 Shell Commands Specification*.

2.1 Objects Managed by EFI-Based Firmware

There are several different types of objects that can be managed through the services provided by EFI. Figure 2-1 below shows the various object types. The most important ones for EFI drivers are the following:

- EFI System Table
- Memory
- Handles
- Images
- Events

Some EFI drivers may need to access environment variables, but most do not. Rarely do EFI drivers require the use of a monotonic counter, watchdog timer, or real-time clock. The EFI System Table is the most important data structure, because it provides access to all the services provided by

EFI and access to all the additional data structures that describe the configuration of the platform. Each of these object types and the services that provide access to them will be introduced in the following sections.

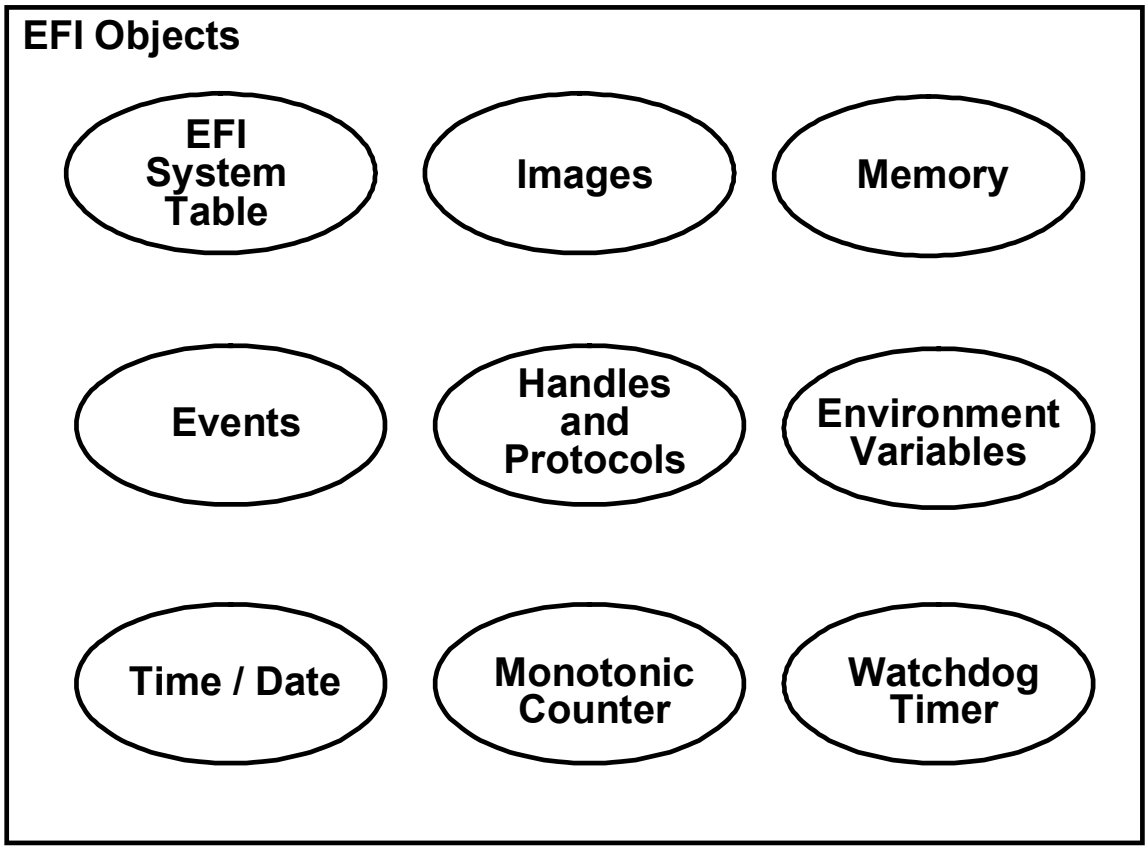

Figure 2-1. Object Managed by EFI-Based Firmware

2.2 EFI System Table

The *EFI System Table* is the most important data structure in EFI. From this one data structure, an EFI executable image can gain access to system configuration information and a rich collection of EFI services. These EFI services include the following:

- EFI Boot Services
- EFI Runtime Services
- Protocol services

The EFI Boot Services and EFI Runtime Services are accessed through the EFI Boot Services Table and the EFI Runtime Services Table, which are two of the data fields in the EFI System Table. The number and type of services that are available from these two tables is fixed for each revision of the *EFI Specification*. The EFI Boot Services and EFI Runtime Services are defined in the *EFI 1.10 Specification*, and the common uses of theses services by EFI drivers are presented in chapter 4 of the *EFI 1.10 Specification.*

Protocol services are groups of related functions and data fields that are named by a Globally Unique Identifier (GUID; see Appendix A of the *EFI 1.10 Specification*). Protocol services are typically used to provide software abstractions for devices such as consoles, disks, and networks. They can also be used to extend the number of generic services that are available in the platform. Protocols are the basic building blocks that allow the functionality of EFI firmware to be extended over time. The *EFI 1.10 Specification* defines over 30 different protocols, and various implementations of EFI firmware and EFI drivers may produce additional protocols to extend the functionality of a platform.

2.3 Handle Database

The *handle database* is composed of objects called handles and protocols. *Handles* are a collection of one or more protocols, and *protocols* are data structures that are named by a GUID. The data structure for a protocol may be empty, may contain data fields, may contain services, or may contain both services and data fields. During EFI initialization, the system firmware, EFI drivers, and EFI applications will create handles and attach one or more protocols to the handles. Information in the handle database is “global” and can be accessed by any executable EFI image.

The handle database is the central repository for the objects that are maintained by EFI-based firmware. The handle database is a list of EFI handles, and each EFI handle is identified by a unique handle number that is maintained by the system firmware. A handle number provides a database “key” to an entry in the handle database. Each entry in the handle database is a collection of one or more protocols. The types of protocols, named by GUID, that are attached to an EFI handle determine the handle type. An EFI handle may represent components such as the following:

- Executable images such as EFI drivers and EFI applications
- Devices such as network controllers and hard drive partitions
- EFI services such as EFI Decompression and the EBC Virtual Machine

Figure 2-2 below shows a portion of the handle database. In addition to the handles and protocols, a list of objects is associated with each protocol. This list is used to track which agents are consuming which protocols. This information is critical to the operation of EFI drivers, because this information is what allows EFI drivers to be safely loaded, started, stopped, and unloaded without any resource conflicts.

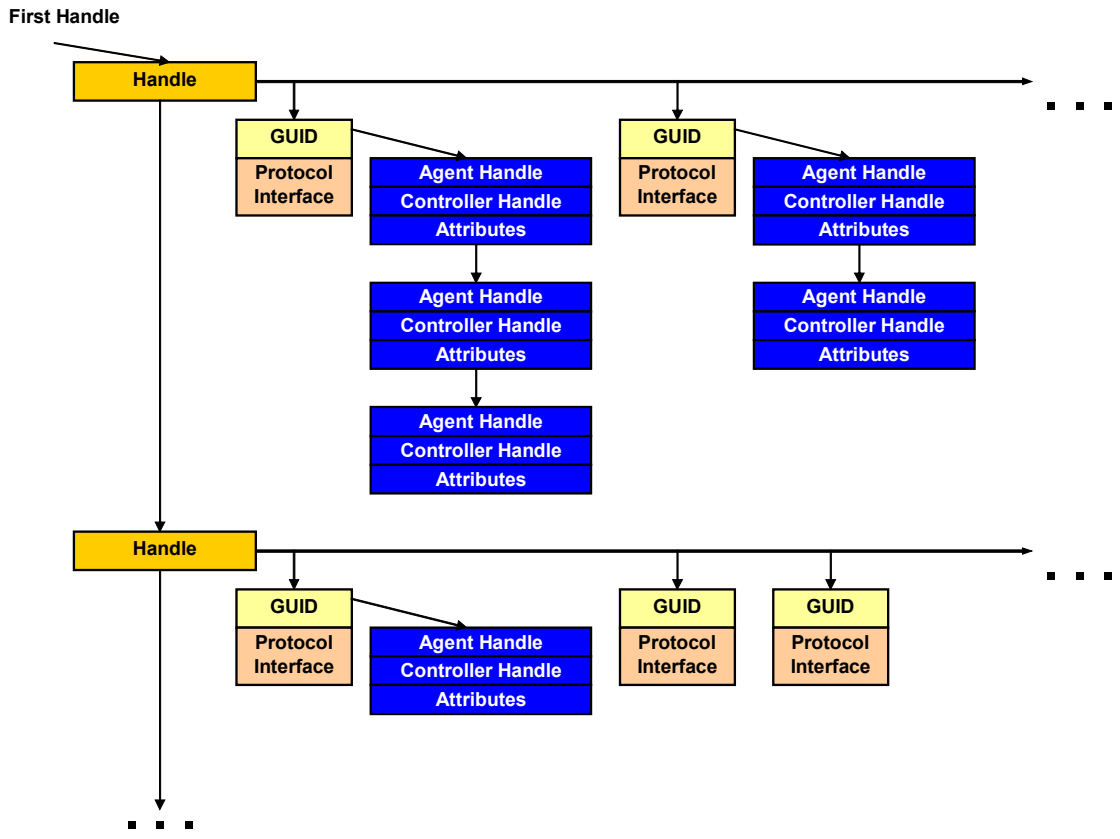

Figure 2-2. Handle Database

Figure 2-3 below shows the different types of handles that may be present in the handle database and the relationships between the various handle types. The handle-related terms introduced here are used throughout the document. All handles reside in the same handle database and the types of protocols that are associated with each handle differentiate the handle type.

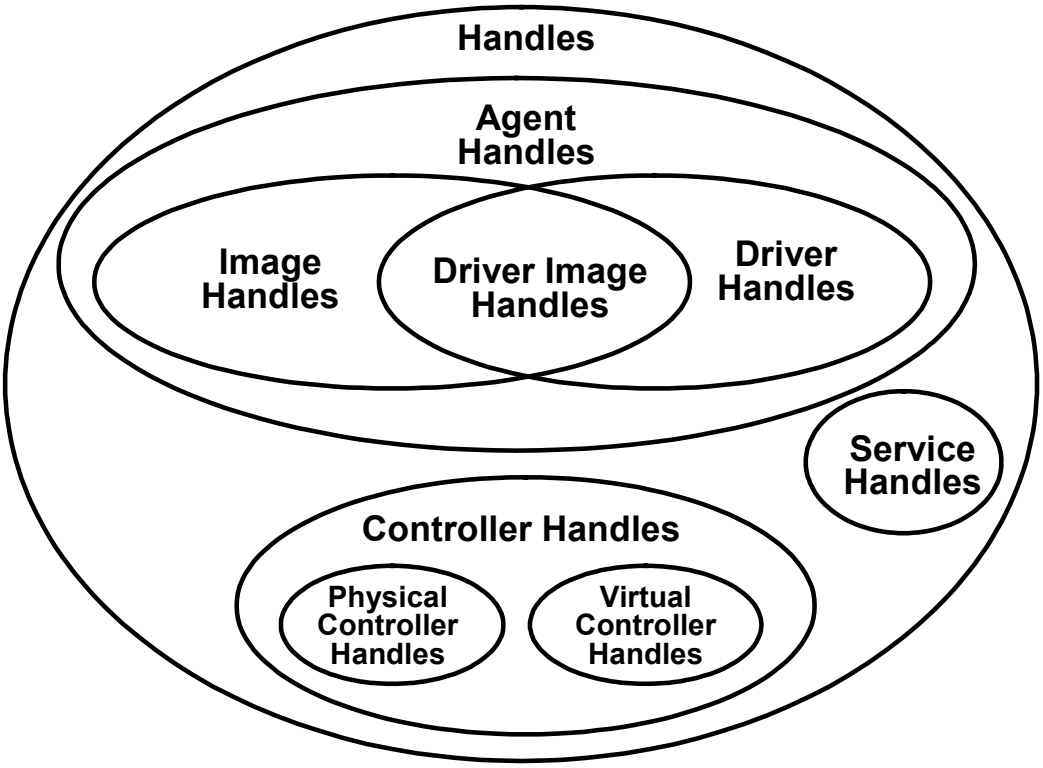

Figure 2-3. Handle Types

Table 2-1 describes the types of handles that are shown in Figure 2-3.

Table 2-1. Description of Handle Types

Type of Handle	Description
Agent handle	This term is used by some of the EFI Driver Model–related services in the *EFI 1.10 Specification*. An *agent* is an EFI component that can consume a protocol in the handle database. An *agent handle* is a general term that can represent an image handle, a driver handle, or a driver image handle.
Image handle	Supports the Loaded Image Protocol. See chapter 7 of the *EFI 1.10 Specification*.
Driver handle	Supports the Driver Binding Protocol. May optionally support the Driver Configuration Protocol, the Driver Diagnostics Protocol, and the Component Name Protocol. *DIG64* requires these optional protocols for Itanium-based platforms. Other platform specifications may or may not require these protocols. See chapter 9 of the *EFI 1.10 Specification*.

continued

Table 2-1. Description of Handle Types (continued)

Type of Handle	Description
Driver image handle	The intersection of image handles and driver handles. Supports both the Loaded Image Protocol and the Driver Binding Protocol. May optionally support the Driver Configuration Protocol, the Driver Diagnostics Protocol, and the Component Name Protocol. *DIG64* requires these optional protocols for Itanium-based platforms. Other platform specifications may or may not require these protocols. See chapter 9 of the *EFI 1.10 Specification*.
Controller handle	If the handle represents a physical device, then it must support the Device Path Protocol. If the handle represents a virtual device, then it must not support the Device Path Protocol. In addition, a device handle must support one or more additional I/O protocols that are used to abstract access to that device. The list of I/O protocols that are defined in the *EFI 1.10 Specification* include the following: • **Console Services:** Simple Input Protocol, Simple Text Output Protocol, UGA Draw Protocol, UGA I/O Protocol, Simple Pointer Protocol, Serial I/O Protocol, and Debug Port Protocol • **Bootable Image Services:** Block I/O Protocol, Disk I/O Protocol, Simple File System Protocol, and Load File Protocol • **Network Services:** Network Interface Identifier Protocol, Simple Network Protocol, and PXE Base Code Protocol • **PCI Services:** PCI Root Bridge I/O Protocol, PCI I/O Protocol, and Device I/O Protocol • **USB Services:** USB Host Controller Protocol, USB I/O Protocol • **SCSI Services:** SCSI Pass Thru Protocol
Device handle	Term is used interchangeably with *controller handle*.
Bus controller handle	Managed by a bus driver or a hybrid driver that produces child handles. The term "bus" does not necessarily match the hardware topology. The term "bus" in this document is used from the software perspective, and the production of the software construct—which is called a *child handle*—is the only distinction between a controller handle and a bus controller handle.
Child handle	A type of controller handle that is created by a bus driver or a hybrid driver. See section 2.6 for the definition of bus driver and hybrid driver. The distinction between a *child handle* and a *controller handle* depends on the perspective of the driver that is using the handle. A handle would be a child handle from a bus driver's perspective, and that same handle may be a controller handle from a device driver's perspective.
Physical controller handle	A controller handle that represents a physical device that must support the Device Path Protocol. See chapter 8 of the *EFI 1.10 Specification.*
Virtual controller handle	A controller handle that represents a virtual device and does not support the Device Path Protocol.

continued

Table 2-1. Description of Handle Types (continued)

Type of Handle	Description
Service handle	Does not support the Loaded Image Protocol, the Driver Binding Protocol, or the Device Path Protocol. Instead, it supports the only instance of a specific protocol in the entire handle database. This protocol provides services that may be used by other EFI applications or EFI drivers. The list of service protocols that are defined in the *EFI 1.10 Specification* include the Platform Driver Override Protocol, Unicode Collation Protocol, Boot Integrity Services Protocol, Debug Support Protocol, Decompress Protocol, and EFI Byte Code Protocol.

2.4 Protocols

The extensible nature of EFI is built, to a large degree, around protocols. EFI drivers are sometimes confused with EFI protocols. Although they are closely related, they are distinctly different. An EFI driver is an executable EFI image that installs a variety of protocols of various handles to accomplish its job.

EFI protocols are a block of function pointers and data structures or APIs that have been defined by a specification. At a minimum, the specification will define a GUID. This number is the protocol's real name and will be used to find this protocol in the handle database. The protocol also typically includes a set of procedures and/or data structures (called the *protocol interface structure*). The following is an example of a protocol definition from section 9.6 of the *EFI 1.10 Specification*. Notice that it defines two function definitions and one data field.

GUID

```
#define EFI_COMPONENT_NAME_PROTOCOL_GUID \
  { 0x107a772c,0xd5e1,0x11d4,0x9a,0x46,0x0,0x90,0x27,0x3f,0xc1,0x4d }
```

Protocol Interface Structure

```
typedef struct _EFI_COMPONENT_NAME_PROTOCOL {
  EFI_COMPONENT_NAME_GET_DRIVER_NAME       GetDriverName;
  EFI_COMPONENT_NAME_GET_CONTROLLER_NAME   GetControllerName;
  CHAR8                                    *SupportedLanguages;
} EFI_COMPONENT_NAME_PROTOCOL;
```

Figure 2-4 below shows a single handle and protocol from the handle database that is produced by an EFI driver. The protocol is composed of a GUID and a protocol interface structure. Many times, the EFI driver that produces a protocol interface will maintain additional private data fields. The protocol interface structure itself simply contains pointers to the protocol function. The protocol functions are actually contained within the EFI driver. An EFI driver may produce one protocol or many protocols depending on the driver's complexity.

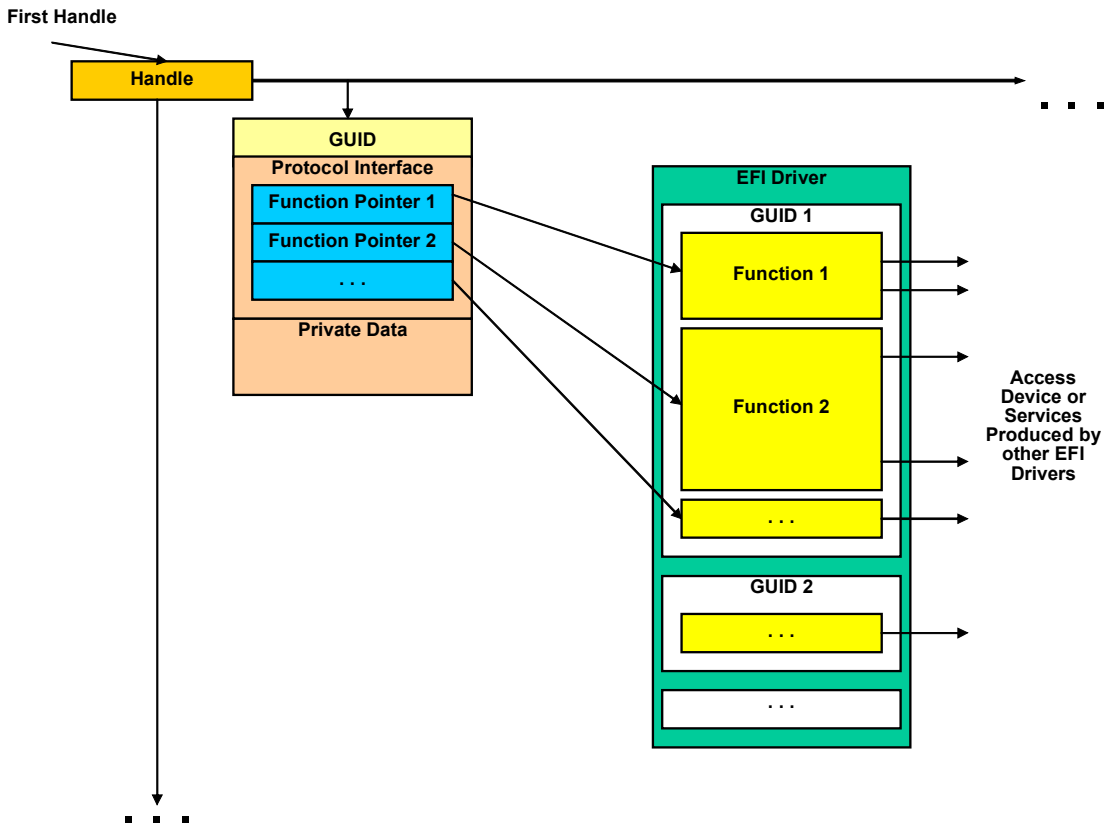

Figure 2-4. Construction of a Protocol

Not all protocols are defined in the *EFI 1.10 Specification*. The *EFI 1.10 Sample Implementation* includes many protocols that are not part of the *EFI 1.10 Specification*. These protocols are necessary to provide all of the functionality in a particular implementation, but they are not defined in the *EFI 1.10 Specification* because they do not present an external interface that is required to support booting an OS or writing an EFI driver. The creation of new protocols is how EFI-based systems can be extended over time as new devices, buses, and technologies are introduced.

The following are a few examples of protocols in the *EFI 1.10 Sample Implementation* that are not part of the *EFI 1.10 Specification:*

- `Varstore`
- `ConIn`
- `ConOut`
- `StdErr`
- `PrimaryConIn`
- `VgaMiniPort`
- `UsbAtapi`

The EFI Application Toolkit also contains a number of EFI protocols, such as the following, that may be found on some platforms:

- PPP Deamon
- Ramdisk
- TCP/IP

The OS loader and drivers should not depend on these types of protocols because they are not guaranteed to be present in every EFI-compliant system. OS loaders and drivers should depend only on protocols that are defined in the *EFI 1.10 Specification* and protocols that are required by platform design guides such as *DIG64*.

The extensible nature of EFI allows each platform to design and add its own special protocols. These protocols can be used to expand that capabilities of EFI and provide access to proprietary devices and interfaces in congruity with the rest of the EFI architecture.

Because a protocol is "named" by a GUID, there should be no other protocols with that same identification number. Care must be taken when creating a new protocol to define a new GUID for it. EFI fundamentally assumes that a specific GUID will expose a specific protocol interface. Cutting and pasting an existing GUID or hand modifying an existing GUID creates the opportunity for a duplicate GUID to be introduced. A system containing a duplicate GUID may inadvertently find the new protocol and think that it is another protocol, which will mostly likely crash the system. These types of bugs are also very difficult to root cause. The versions of Microsoft Visual Studio that are referenced by the *EFI 1.10 Sample Implementation's* RELNOTES.DOC include a **`GUIDGEN`** tool that can quickly generate a GUID definition that is suitable for including in the code.

2.4.1 Working with Protocols

Any EFI code can operate with protocols during boot time. However, after **`ExitBootServices()`** is called, the handle database is no longer available. There are several EFI boot time services for working with EFI protocols. Section 5.3 of the *EFI 1.10 Specification* discusses these Protocol Handler Services. These services are described in more detail in chapter 4 of this document.

2.4.2 Multiple Protocol Instances

A handle may have many protocols attached to it. However, it may have only one protocol of each type. In other words, a handle may not have more than one of the exact same protocol. Otherwise, it would make requests for a particular protocol on a handle nondeterministic.

However, drivers may create multiple "instances" of a particular protocol and attach each instance to a different handle. This scenario is the case with the PCI I/O Protocol, where the PCI bus driver installs a PCI I/O Protocol instance for each PCI device. Each "instance" of the PCI I/O Protocol is configured with data values that are unique to that PCI device, including the location and size of the EFI Option ROM (OpROM) image.

Also, each driver can install customized versions of the same protocol (as long as it is not on the same handle). For example, each EFI driver installs the Component Name Protocol on its driver image handle, yet when the **EFI_COMPONENT_NAME_PROTOCOL.GetDriverName()** function is called, each handle will return the unique name of the driver that owns that image handle. The **EFI_COMPONENT_NAME_PROTOCOL.GetDriverName()** function on the USB bus driver handle will return "USB bus driver" for the English language, but the **EFI_COMPONENT_NAME_PROTOCOL.GetDriverName()** function on the PXE driver handle will return "PXE base code driver."

2.4.3 Tag GUID

A protocol may be nothing more than a GUID. This GUID is also known as a *tag GUID*. Such a protocol can still serve very useful purposes such as marking a device handle as special in some way or allowing other EFI images to easily find the device handle by querying the system for the device handles with that protocol GUID attached.

The *EFI 1.10 Sample Implementation* uses the **HOT_PLUG_DEVICE_GUID** in this way to mark device handles that represent devices from a hot-plug bus such as USB.

2.5 EFI Images

This section describes the different types of EFI images. All EFI images contain a PE/COFF header that defines the format of the executable code (see the *Microsoft Portable Executable and Common Object File Format Specification* for more information). The code can be for IA-32, the Itanium processor, or EFI Byte Code. The header will define the processor type and the image type. Section 2.1.1 of the *EFI 1.10 Specification* defines the three processor types and the following three image types:

- EFI applications
- EFI Boot Service drivers
- EFI Runtime drivers

EFI images are loaded and relocated into memory with the Boot Service **gBS->LoadImage()**. There are several supported storage locations for EFI images, including the following:

- Expansion ROMs on a PCI card
- System ROM or system flash
- A media device such as a hard disk, floppy, CD-ROM, or DVD
- LAN boot server

In general, EFI images are not compiled and linked at a specific address. Instead, they are compiled and linked such that relocation fix-ups are included in the EFI image, which allows the EFI image to be placed anywhere in system memory. The Boot Service **gBS->LoadImage()** does the following:

- Allocates memory for the image being loaded
- Automatically applies the relocation fix-ups to the image
- Creates a new image handle in the handle database, which installs an instance of the **EFI_LOADED_IMAGE_PROTOCOL**

This instance of the **EFI_LOADED_IMAGE_PROTOCOL** contains information about the EFI image that was loaded. Because this information is published in the handle database, it is available to all EFI components.

After an EFI image is loaded with **gBS->LoadImage()**, it can be started with a call to **gBS->StartImage()**. The header for an EFI image contains the address of the entry point that is called by **gBS->StartImage()**. The entry point always receives the following two parameters:

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

These two items allow the EFI image to do the following:

- Access all of the EFI services that are available in the platform.
- Retrieve information about where the EFI image was loaded from and where in memory the image was placed.

The operations that are performed by the EFI image in its entry point vary depending on the type of EFI image. Figure 2-5 below shows the various EFI image types and the relationships between the different levels of images.

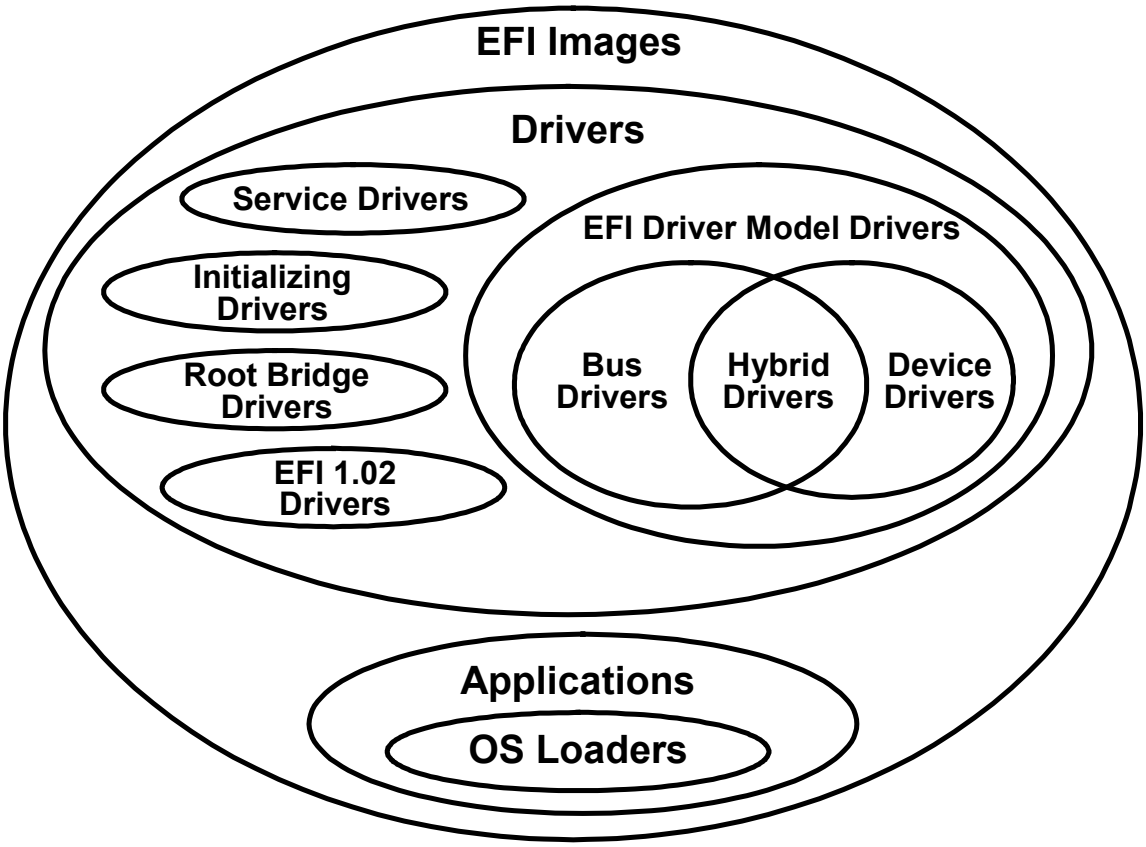

Figure 2-5. Image Types

Table 2-2 describes the types of images that are shown in Figure 2-5.

Table 2-2. Description of Image Types

Type of Image	Description
Application	An EFI image of type **EFI_IMAGE_SUBSYSTEM_EFI_APPLICATION**. This image is executed and automatically unloaded when the image exits or returns from its entry point.
OS loader	A special type of application that normally does not return or exit. Instead, it calls the EFI Boot Service **gBS->ExitBootServices()** to transfer control of the platform from the firmware to an operating system.
Driver	An EFI image of type **EFI_IMAGE_SUBSYSTEM_BOOT_SERVICE_DRIVER** or **EFI_IMAGE_SUBSYSTEM_RUNTIME_DRIVER**. If this image returns **EFI_SUCCESS**, then the image is not unloaded. If the image returns an error code other than **EFI_SUCCESS**, then the image is automatically unloaded from system memory. The ability to stay resident in system memory is what differentiates a driver from an application. Because drivers can stay resident in memory, they can provide services to other drivers, applications, or an operating system. Only the services produced by runtime drivers are allowed to persist past **gBS->ExitBootServices()**.
Service driver	A driver that produces one or more protocols on one or more new service handles and returns **EFI_SUCESS** from its entry point.
Initializing driver	A driver that does not create any handles and does not add any protocols to the handle database. Instead, this type of driver performs some initialization operations and returns an error code so the driver is unloaded from system memory.
Root bridge driver	A driver that creates one or physical controller handles that contain a Device Path Protocol and a protocol that is a software abstraction for the I/O services provided by a root bus produced by a core chipset. The most common root bridge driver is one that creates handles for the PCI root bridges in the platform that support the Device Path Protocol and the PCI Root Bridge I/O Protocol.
EFI 1.02 driver	A driver that follows the *EFI 1.02 Specification*. This type of driver does not use the EFI Driver Model. These types of drivers are not discussed in detail in this document. Instead, this document presents recommendations on converting EFI 1.02 drivers to drivers that follow the EFI Driver Model.

continued

Table 2-2. Description of Image Types (continued)

Type of Image	Description
EFI Driver Model driver	A driver that follows the EFI Driver Model that is described in detail in the *EFI 1.10 Specification*. This type of driver is fundamentally different from service drivers, initializing drivers, root bridge drivers, and EFI 1.02 drivers because a driver that follows the EFI Driver Model is not allowed to touch hardware or produce device-related services in the driver entry point. Instead, the driver entry point of a driver that follows the EFI Driver Model is allowed only to register a set of services that allow the driver to be started and stopped at a later point in the system initialization process.
Device driver	A driver that follows the EFI Driver Model. This type of driver produces one or more driver handles or driver image handles by installing one or more instances of the Driver Binding Protocol into the handle database. This type of driver does not create any child handles when the `Start()` service of the Driver Binding Protocol is called. Instead, it only adds additional I/O protocols to existing controller handles.
Bus driver	A driver that follows the EFI Driver Model. This type of driver produces one or more driver handles or driver image handles by installing one or more instances of the Driver Binding Protocol in the handle database. This type of driver creates new child handles when the `Start()` service of the Driver Binding Protocol is called. It also adds I/O protocols to these newly created child handles.
Hybrid driver	A driver that follows the EFI Driver Model and shares characteristics with both device drivers and bus drivers. This distinction means that the `Start()` service of the Driver Binding Protocol will add I/O protocols to existing handles and it will create child handles.

2.5.1 Applications

An EFI application will start execution at its entry point and then execute until it returns from its entry point or it calls the `Exit()` boot service function. When it is done, the image is unloaded from memory. It does not stay resident in memory. Some examples of common EFI applications include the following:

- EFI Shell
- EFI Shell commands
- Flash utilities
- Diagnostic utilities

It is perfectly acceptable to invoke EFI applications from inside other EFI applications.

2.5.2 OS Loader

The *EFI 1.10 Specification* details a special type of EFI application called an *OS boot loader*. It is an EFI application that calls **`ExitBootServices()`**. **`ExitBootServices()`** is called when the OS loader has set up enough of the OS infrastructure that it is ready to assume ownership of the system resources. At **`ExitBootServices()`**, the EFI core will free all of its boot time services and drivers, leaving only the runtime services and drivers.

2.5.3 Drivers

EFI drivers are different from EFI applications in that unless there is an error returned from the driver's entry point, the driver stays resident in memory. The EFI core firmware, the boot manager, or other EFI applications may load drivers.

2.5.3.1 EFI 1.02 Drivers

There are several types of EFI drivers. In EFI 1.02, drivers were constructed without a defined driver model. The *EFI 1.10 Specification* provides a driver model that replaces the way drivers were built in EFI 1.02 but that still maintains backward compatibility with EFI 1.02 drivers. EFI 1.02 immediately started the driver inside the entry point. This method meant that the driver would immediately search for supported devices, install the necessary I/O protocols, and start the timers that were needed to poll the devices. However, this method did not allow the system to have control over the driver loading and connection policies, so the EFI Driver Model was introduced in the *EFI 1.10 Specification* to resolve these issues. See section 1.6.1 of the *EFI 1.10 Specification* for details.

The Floating Point Software Assist (FPSWA) driver is a common example of an EFI 1.02 driver; other EFI 1.02 drivers can be found in the EFI Application Toolkit 1.02.12.38. It is strongly recommended that EFI 1.02 drivers be converted to EFI 1.10 drivers that follow the EFI Driver Model.

2.5.3.2 Boot Service and Runtime Drivers

Boot time drivers are loaded into memory marked as **`EfiBootServicesCode`**, and they allocate their data structures from memory marked as **`EfiBootServicesData`**. These memory types are converted to available memory after **`gBS->ExitBootServices()`** is called.

Runtime drivers are loaded in memory marked as **`EfiRuntimeServicesCode`**, and they allocate their data structures from memory marked as **`EfiRuntimeServicesData`**. These types of memory are preserved after **`gBS->ExitBootServices()`** is called. This preservation allows runtime driver to provide services to an operating system while the operating system is running. Runtime drivers must publish an alternative calling mechanism, because the EFI handle database does not persist into OS runtime. The most common examples of EFI runtime drivers are the Floating Point Software Assist driver (**`FPSWA.efi`**) and the network Universal Network Driver Interface (UNDI) driver. Other than these examples, runtime drivers are not very common and will not be discussed in detail. In addition, the implementation and validation of runtime drivers is much more difficult than boot service drivers because EFI supports the translation of runtime services and runtime drivers from a physical addressing mode to a virtual addressing mode.

2.6 Events and Task Priority Levels

Events are another type of object that is managed through EFI services. They can be created and destroyed and are either in the waiting state or the signaled state. An EFI image can do any of the following:

- Create an event.
- Destroy an event.
- Check to see if an event is in the signaled state.
- Wait for an event to be in the signaled state.
- Request that an event be moved from the waiting state to the signaled state.

Because EFI does not support interrupts, it can present a challenge to driver writers who are accustomed to an interrupt-driven driver model. Instead, EFI supports polled drivers. The most common use of events by an EFI driver is the use of timer events that allows drivers to periodically poll a device. Figure 2-6 below shows the different types of events that are supported in EFI and the relationships between those events.

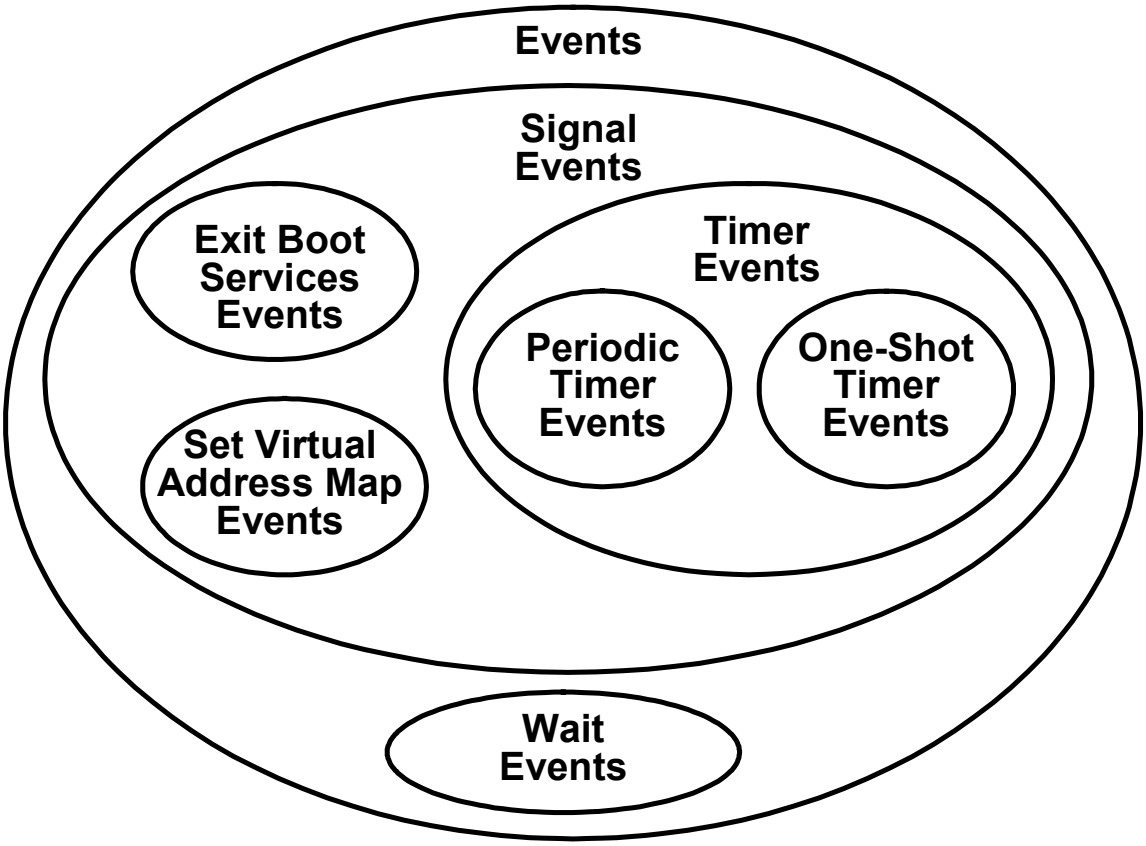

Figure 2-6. Event Types

Table 2-3 describes the types of events that are shown in Figure 2-6.

Table 2-3. Description of Event Types

Type of Events	Description
Wait event	An event whose notification function is executed whenever the event is checked or waited upon.
Signal event	An event whose notification function is scheduled for execution whenever the event goes from the waiting state to the signaled state.
Exit Boot Services event	A special type of signal event that is moved from the waiting state to the signaled state when the EFI Boot Service **`ExitBootServices()`** is called. This call is the point in time when ownership of the platform is transferred from the firmware to an operating system. The event's notification function is scheduled for execution when **`ExitBootServices()`** is called.
Set Virtual Address Map event	A special type of signal event that is moved from the waiting state to the signaled state when the EFI Runtime Service **`SetVirtualAddressMap()`** is called. This call is the point in time when the operating system is making a request for the runtime components of EFI to be converted from a physical addressing mode to a virtual addressing mode. The operating system provides the map of virtual addresses to use. The event's notification function is scheduled for execution when **`SetVirtualAddressMap()`** is called.
Timer event	A type of signal event that is moved from the waiting state to the signaled state when at least a specified amount of time has elapsed. Both periodic and one-shot timers are supported. The event's notification function is scheduled for execution when a specific amount of time has elapsed.
Periodic timer event	A type of timer event that is moved from the waiting state to the signaled state at a specified frequency. The event's notification function is scheduled for execution when a specific amount of time has elapsed.
One-shot timer event	A type of timer event that is moved from the waiting state to the signaled state after the specified timer period has elapsed. The event's notification function is scheduled for execution when a specific amount of time has elapsed.

The following three elements are associated with every event:

- The Task Priority Level (TPL) of the event
- A notification function
- A notification context

The notification function for a wait event is executed when the state of the event is checked or when the event is being waited upon. The notification function of a signal event is executed whenever the event transitions from the waiting state to the signaled state. The notification context is passed into the notification function each time the notification function is executed. The TPL is the priority at which the notification function is executed. Table 2-4 lists the four TPL levels that are defined in EFI.

Table 2-4. Task Priority Levels Defined in EFI

Task Priority Level	Description
TPL_APPLICATION	The priority level at which EFI images are executed.
TPL_CALLBACK	The priority level for most notification functions.
TPL_NOTIFY	The priority level at which most I/O operations are performed.
TPL_HIGH_LEVEL	The priority level for the one timer interrupt supported in EFI.

TPLs serve the following two purposes:

- To define the priority in which notification functions are executed
- To create locks

In the first purpose, this mechanism is used only when more than one event is in the signaled state at the same time. In these cases, the notification function that has been registered with the higher priority will be executed first. Also, notification functions at higher priorities can interrupt the execution of notification functions executing at a lower priority.

In the second purpose, creating locks, it is possible for the code in normal context and the code in interrupt context to access the same data structure, because EFI does support a single timer interrupt. This access can cause issues if the updates to a shared data structure are not atomic. An EFI application or EFI driver that wants to guarantee exclusive access to a shared data structure can temporarily raise the task priority level to prevent simultaneous access from both normal context and interrupt context. A lock can be created by temporarily raising the task priority level to **TPL_HIGH_LEVEL**. This level will even block the one timer interrupt, but care must be taken to minimize the amount of time that the system is at **TPL_HIGH_LEVEL**, because all timer-based events are blocked during this time and any driver that requires periodic access to a device will be prevented from accessing its device.

2.7 EFI Device Paths

EFI defines a Device Path Protocol that is attached to device handles in the handle database to help operating systems and their loaders identify the hardware that a device handle represents. The Device Path Protocol provides a unique name for each physical device in a system. The collection of Device Path Protocols for the physical devices managed by EFI-based firmware is called a name space. Modern operating systems tend to use ACPI and industry standard buses to produce a name space while the operating system is running. However, the ACPI name space is difficult to parse, and it would greatly increase the size and complexity of system firmware to carry an ACPI name space parser. Instead, EFI uses aspects of the ACPI name space that do not require an ACPI name space parser. This compromise keeps the size and complexity of system firmware to a minimum and provides a way for the operating system to create a mapping from EFI device paths to the operating system's name space.

A *device path* is a data structure that is composed of one or more device path nodes. Every device path node contains a standard header that includes the node's type, subtype, and length. This standard header allows a parser of a device path to hop from one node to the next without having to understand every type of node that may be present in the system. Example 2-1 below shows the declaration of the standard header for an EFI device path. The device or bus-specific node data follows the `Length` field. The declaration of the PCI device path node is also shown in the example below. It contains the PCI-device-specific fields `Function` and `Device`.

```
//
// Generic device path node header
//
typedef struct
  UINT8  Type;
  UINT8  SubType;
  UINT8  Length[2];
} EFI_DEVICE_PATH_PROTOCOL;

//
// PCI Device Path Node that includes a header and PCI-specific fields
//
typedef struct _PCI_DEVICE_PATH {
  EFI_DEVICE_PATH_PROTOCOL        Header;
  UINT8                           Function;
  UINT8                           Device;
} PCI_DEVICE_PATH;
```

Example 2-1. EFI Device Path Header

A device path is position independent because it is not allowed to contain any pointer values. This independence allows device paths to be easily moved from one location to another and stored in nonvolatile storage. A device path is terminated by a special device path node called an *end device path node*. Table 2-5 lists the types of device path nodes that are defined in section 8.3 of the *EFI 1.10 Specification.*

Table 2-5. Types of Device Path Nodes Defined in EFI 1.10 Specification

Type of Device Path Nodes	Description
Hardware device path node	Used to describe devices on industry-standard buses that are directly accessible through processor memory or processor I/O cycles. These devices include memory-mapped devices and devices on PCI buses and PC card buses.
ACPI device path node	Used to describe devices whose enumeration is not described in an industry-standard fashion. This type of device path is used to describe devices such as PCI root bridges and ISA devices. These device path nodes contain HID, CID, and UID fields that must match the HID, CID, and UID values that are present in the platform's ACPI tables.
Messaging device path node	Used to describe devices on industry-standard buses that are not directly accessible through processor memory or processor I/O cycles. These devices are accessed by the processor through one or more hardware bridge devices that translate one industry-standard bus type to another industry-standard bus type. This type of device path is used to describe devices such as SCSI, Fibre Channel, 1394, USB, I2O, InfiniBand*, UARTs, and network agents.
Media device path node	Hard disk, CD-ROM, and file paths in a file system that support multiple directory levels.
BIOS Boot Specification device path node	Used to describe a device that has a type that follows the *BIOS Boot Specification,* such as floppies, hard disks, CD-ROMs, PCMCIA devices, USB devices, network devices, and Bootstrap Entry Vector (BEV) devices. These device path nodes are used only in a platform that supports BIOS INT services.
End device path node	Used to terminate a device path.

Each of the device path node types also supports a vendor-defined node that is the extensibility mechanism for device paths. As new devices, bus types, and technologies are introduced into platforms, new device path nodes may have to be created. The vendor-defined nodes use a GUID to distinguish new device path nodes. When a device path node is created for a new device or bus type, the **GUIDGEN** tool should be used to create a new GUID. Care must be taken in the choice of the data fields used in the definition of a new device path node. As long as a device is not physically moved from one location in a platform to another location, the device path must not change across platform boots or system configuration changes. For example, the PCI device path node only contains a *Device* and a *Function* field. It does not contain a *Bus* field, because the addition of a device with a PCI-to-PCI bridge may modify the bus numbers of other devices in the platform. Instead, the device path for a PCI device is described with one or more PCI device path nodes that describe the path from the PCI root bridge, through zero or more PCI-to-PCI bridges, and finally the target PCI device.

The EFI Shell is able to display a device path on a console as a string. The conversion of device path nodes to printable strings is not architectural. It is a feature that is implemented in the EFI Shell that allows developers to view device paths without having to manually translate hex dumps of the device path node data structures. Example 2-2 shows some example device paths from Itanium and IA-32 platforms. These device paths show standard and extended ACPI device path nodes being used for a PCI root bridge and an ISA floppy controller. PCI device path nodes are

used for PCI-to-PCI bridges, PCI video controllers, PCI IDE controllers, and PCI-to-LPC bridges. Finally, IDE messaging device path nodes are used to describe an IDE hard disk, and media device path nodes are used to describe a partition on an IDE hard disk.

```
//
// PCI Root Bridge #0 using an Extended ACPI Device Path
//
Acpi(HWP0002,PNP0A03,0)

//
// PCI Root Bridge #1 using an Extended ACPI Device Path
//
Acpi(HWP0002,PNP0A03,1)

//
// PCI Root Bridge #0 using a standard ACPI Device Path
//
Acpi(PNP0A03,0)

//
// PCI-to-PCI bridge device directly attached to PCI Root Bridge #0
//
Acpi(PNP0A03,0)/Pci(1E|0)

//
// A video adapter installed in a slot on the other side of a PCI-to-PCI bridge
// that is attached to PCI Root Bridge #0.
//
Acpi(PNP0A03,0)/Pci(1E|0)/Pci(0|0)

//
// A PCI-to-LPC bridge device attached to PCI Root Bridge #0
//
Acpi(PNP0A03,0)/Pci(1F|0)
//
// A 1.44 MB floppy disk controller attached to a PCI-to-LPC bridge device
// attached to PCI Root Bridge #0
//
Acpi(PNP0A03,0)/Pci(1F|0)/Acpi(PNP0604,0)

//
// A PCI IDE controller attached to PCI Root Bridge #0
//
Acpi(PNP0A03,0)/Pci(1F|1)

//
// An IDE hard disk attached to a PCI IDE controller attached to
// PCI Root Bridge #0
//
Acpi(PNP0A03,0)/Pci(1F|1)/Ata(Secondary,Master)

//
// Partition #1 of an IDE hard disk attached to a PCI IDE controller attached to
// PCI Root Bridge #0
//
Acpi(PNP0A03,0)/Pci(1F|1)/Ata(Secondary,Master)/HD(Part1,Sig00000000)
```

Example 2-2. Example Device Paths

2.7.1 How Drivers Use Device Paths

EFI drivers that manage physical devices must be aware of device paths. When possible, EFI drivers treat device paths as opaque data structures. Device drivers do not operate on them at all. Root bridge drivers are required only to produce the device paths for the root bridges, which typically contain only a single ACPI device path node, and bus drivers usually just append a single device path node for a child device to the device path of the parent device. The bus drivers should not parse the contents of the parent device path. Instead, a bus driver appends the one device path node that it is required to understand to the device path of the parent device.

For example, a SCSI driver that produces child handles for the disk devices on the SCSI channel will need to build a device path for each disk device. The device path will be constructed by appending a SCSI device path node to the device path of the SCSI controller itself. The SCSI device path node simply contains the Physical Unit Number (PUN) and Logical Unit Number (LUN) of the SCSI disk device, and the driver for a SCSI host adapter is already tracking that information.

The mechanism described above allows the construction of device paths to be a distributed process. The bus drivers at each level of the system hierarchy are required only to understand the device path nodes for their child devices. Bus drivers understand their local view of the device path, and a group of bus drivers from each level of the system bus hierarchy work together to produce complete device paths for the console and boot devices that are used to install and boot operating systems.

There are a number of EFI Library services that are available to help manage device paths. Please see the *EFI Library Specification* for more details.

2.7.2 Considerations for Itanium® Architecture

Device paths nodes can be any length, which can actually cause problems on Itanium-based platforms where all data that is accessed must be aligned on proper boundaries. A device path node that is not a multiple of 8 bytes in length will cause the device paths nodes that follow it to be unaligned. Care must be taken when using device paths to make sure an alignment fault is not generated. Device paths that are stored in environment variables are packed on purpose to reduce the amount of nonvolatile storage that is consumed by device paths. These device paths can be unpacked, so each device path node is guaranteed to be on an 8-byte boundary. However, there is no standard way to tell if a device path is packed, unpacked, aligned, or unaligned based on the device path contents alone. Instead, consumers of device paths should always take measures to guarantee that an alignment fault is never generated. The basic technique is to copy a device path node from a potentially unaligned device path into a device path node that is known to be aligned. Then, the device path node contents can be examined or updated and potentially copied back to the original device path. Some example of these operations can be found in section 18.1.

2.7.3 Environment Variables

Device paths are also used when environment variables are built and stored in nonvolatile storage. There are a number of environment variables defined in section 3.2 of the *EFI 1.10 Specification*. These variables define the following:

- Console input devices
- Console output devices
- Standard error devices
- The drivers that need to be loaded prior to an OS boot
- The boot selections that the platform supports

The EFI boot manager, EFI utilities, and EFI-compliant operating systems manage these environment variables as operating systems are installed and removed from a platform.

2.8 EFI Driver Model

The *EFI 1.10 Specification* defines the EFI Driver Model. Drivers that follow the EFI Driver Model share the same image characteristics as EFI applications. However, the model allows EFI more control over drivers by separating the loading of drivers into memory from the starting and stopping of drivers. Table 2-6 lists the series of EFI Driver Model–related protocols that are used to accomplish this separation.

Table 2-6. Protocols Used to Separate the Loading and Starting/Stopping of Drivers

Protocol	Description
Driver Binding Protocol	Provides functions for starting and stopping the driver, as well as a function for determining if the driver can manage a particular controller. The EFI Driver Model requires this protocol.
Component Name Protocol	Provides functions for retrieving a human-readable name of a driver and the controllers that a driver is managing. While the *EFI 1.10 Specification* lists this protocol as optional, the *Developer's Interface Guide for 64-bit Intel Architecture-based Servers* (hereafter referred to as "DIG64 specification" or just "*DIG64*") lists this protocol as required for Itanium-based platforms.
Driver Diagnostics Protocol	Provides functions for executing diagnostic functions on the devices that a driver is managing. While the *EFI 1.10 Specification* lists this protocol as optional, *DIG64* lists this protocol as required for Itanium-based platforms.
Driver Configuration Protocol	Provides functions that allow the user to configure the devices that a driver is managing. It also allows a device to be placed in its default configuration. While the *EFI 1.10 Specification* lists this protocol as optional, *DIG64* lists this protocol as required for Itanium-based platforms.

The new protocols are registered on the driver's image handle. In the EFI Driver Model, the main goal of the driver's entry point is to install theses protocols and exit successfully. At a later point in the system initialization, EFI can use these protocol functions to operate the driver. A more complex driver may produce more than one instance of the **EFI_DRIVER_BINDING_PROTOCOL**. In this case, the additional instances of the **EFI_DRIVER_BINDING_PROTOCOL** will be installed on new handles. These new handles may also optionally support the Component Name Protocol, Driver Configuration Protocol, and Driver Diagnostics Protocol.

The EFI Driver Model follows the organization of physical/electrical architecture by defining the following two basic types of EFI boot time drivers:

- Device drivers
- Bus drivers

A driver that has characteristics of both a device driver and a bus driver is known as a *hybrid driver*.

Device drivers and bus drivers are distinguished by the operations that they perform in the **Start()** and **Stop()** services of the **EFI_DRIVER_BINDING_PROTOCOL**. By walking through the process of connecting a driver to a device, the roles and relationships of the bus drivers and device drivers will become evident; the following sections discuss these two driver types.

2.8.1 Device Driver

The **Start()** service a device driver will install the protocol(s) directly onto the controller handle that was passed into the **Start()** service. The protocol(s) installed by the device driver uses the I/O services that are provided by the bus I/O protocol that is installed on the controller handle. For example, a device driver for a PCI controller would use the services of the PCI I/O Protocol, and a device driver for a USB device would use the service of the USB I/O Protocol. This process is called "consuming the bus I/O abstraction." The following are the main objectives of the device driver:

- Initialize the controller.
- Install an I/O protocol on the device that can be used by the EFI system firmware to boot the operating system.

It does not make sense to write device drivers for devices that cannot be used to boot a platform. Table 2-7 lists the standard I/O protocols that the *EFI 1.10 Specification* defines for different classes of devices.

Table 2-7. I/O Protocols Used for Different Device Classes

Class of Device	Protocol Used	Defined in this section in *EFI 1.10 Specification*
Media	**BLOCK_IO_PROTOCOL**	11.6
LAN	**NETWORK_INTERFACE_IDENTIFIER_PROTOCOL** (also known as UNDI)	15.2
Console output	**SIMPLE_TEXT_OUTPUT_PROTOCOL**	10.3

continued

Table 2-7. I/O Protocols Used for Different Device Classes (continued)

Class of Device	Protocol Used	Defined in this section in *EFI 1.10 Specification*
Console input	**SIMPLE_INPUT_PROTOCOL**	10.2
Root bus	**EFI_PCI_ROOT_BRIDGE_IO_PROTOCOL** [Note]	12.2

Note Provided by the system firmware. Drivers that produce this protocol are a special case because it is the first device in the handle database (called a "root" device). Because of this, it does not follow all the normal EFI Driver Model rules and is not a good example driver to follow when making your own device drivers.

The fundamental EFI definition of a device driver is that it does not create any child handles. This capability distinguishes a device driver from a bus driver. This definition has the potential to create some confusion because it is often reasonable to write a driver that creates child handles. This creation will make the driver a bus driver by definition, but it may not be managing a hardware bus in the classical sense (such as a PCI, SCSI, USB, or Fibre Channel bus).

Just because a device driver does not create child handles does not mean that the device that a device driver is managing will never be a "parent." The protocol(s) produced by a device driver on a controller handle may be consumed by a bus driver that produces child handles. In this case, the controller handle that is managed by a device driver is a parent controller. This scenario happens quite often. For example, the **EFI_USB_HC_PROTOCOL** is produced by a device driver called the *USB host controller driver*. This protocol is consumed by the bus driver that is called the *USB bus driver* and that creates child handles that contain the **USB_IO_PROTOCOL**. The USB host controller driver that produced the **EFI_USB_HC_PROTOCOL** has no knowledge of the child handles that are produced by the USB bus driver.

2.8.2 Bus Driver

A bus driver is nearly identical to a device driver except that it creates child handles. This ability leads to several added features and responsibilities for a bus driver that will be addressed in detail throughout this document. For example, device drivers do not need to concern themselves with searching the bus.

Just like a device driver, the **Start()** function of a bus driver will consume the parent bus I/O abstraction(s) and produce new I/O abstractions in the form of protocols. For example, the PCI bus driver consumes the services of the **EFI_PCI_ROOT_BRIDGE_IO_PROTOCOL** and uses these services to scan a PCI bus for PCI controllers. Each time that a PCI controller is found, a child handle is created and the **EFI_PCI_IO_PROTOCOL** is installed on the child handle. The services of the **EFI_PCI_IO_PROTOCOL** are implemented using the services of the **EFI_PCI_ROOT_BRIDGE_IO_PROTOCOL**. As a second example, the USB bus driver uses the services of the **EFI_USB_HC_PROTOCOL** to discover and create child handles that support the **EFI_USB_IO_PROTOCOL** for each USB device on the USB bus. The services of the **EFI_USB_IO_PROTOCOL** are implemented using the services of the **EFI_USB_HC_PROTOCOL**.

The following are the main objectives of the bus driver:

- Initialize the bus controller.

- Determine how many children to create. For example, the PCI bus driver may discover and enumerate all PCI devices on the bus or only a single PCI device that is being used to boot. How a bus driver handles this step creates a basic subdivision in the types of bus drivers. A bus driver can do one of the following:
  - Create handles for all child controllers on the first call to **`Start()`**.
  - Allow the handles for the child controllers to be created across multiple calls to **`Start()`**.

  This second type of bus driver is very useful because it reduces the platform boot time. It allows a few child handles or even a single child handle to be created. On buses that take a long time to enumerate their children (for example, SCSI and Fibre Channel), multiple calls to **`Start()`** can save a large amount of time when booting a platform.
- Allocate resources and create a child handle in the EFI handle database for one or more child controllers.
- Install an I/O protocol on the child handle that abstracts the I/O operations that the controller supports (such as the PCI I/O Protocol or the USB I/O Protocol).
- If the child handle represents a physical device, then install a Device Path Protocol (see chapter 8 in the *EFI 1.10 Specification*).
- Load drivers from option ROMs if present. The PCI bus driver is the only bus driver that loads from option ROMs so far.

Some common examples of EFI bus drivers include the following:

- PCI bus driver
- USB bus driver
- SCSI bus drivers

Because bus drivers are defined as drivers that produce child handles, there are some other drivers, such as the following, that unexpectedly qualify as bus drivers:

- **Serial driver:** Creates a child handle and extends the Device Path Protocol to include a **`Uart()`** messaging device path node.
- **LAN driver:** Creates a child handle and extends the Device Path Protocol to include a **`Mac()`** address messaging device path node.

2.9 Simplified Connection Process

All EFI drivers that adhere to the EFI Driver Model follow the same basic procedure. When the driver is loaded, it will install a Driver Binding Protocol on the image handle from which it was loaded. It may also update a pointer to the **`EFI_LOADED_IMAGE`** Protocol's **`Unload()`** service. The driver will then exit from the entry point, leaving the code resident in system memory.

The Driver Binding Protocol provides a version number and the following three services:

- **`Supported()`**
- **`Start()`**
- **`Stop()`**

These services are available on the driver's image handle after the entry point is exited. Later on when the system is "connecting" drivers to devices in the system, the driver's Driver Binding Protocol **Supported()** service is called. The **Supported()** service is passed a controller handle. Quickly, the **Supported()** function will examine the controller handle to see if it represents a device that the driver knows how to manage. If so, it will return **EFI_SUCCESS**. The system will then start the driver by calling the driver's **Start()** service, passing in the supported controller handle. The driver can later be disconnected from a controller handle by calling the **Stop()** service.

This section will walk the reader through the connection process that a platform may undergo to connect the devices in the system with the drivers that are available in the system. This process will appear complex at first, but as the process continues, it will become evident that the same simple procedures are used to accomplish the complex task. This description does not go into all the details of the connection process but explains enough that the role of various drivers in the connection process can be clearly understood. This knowledge is fundamental to designing new EFI drivers.

The EFI Boot Service **ConnectController()** clearly demonstrates the power of the EFI Driver Model. The EFI Shell command **connect** directly exposes much of the functionality of this boot service and provides a convenient way to explore the flexibility and control offered by **ConnectController()**.

2.9.1 ConnectController()

By passing the handle of a specific controller into **ConnectController()**, EFI will follow a specific process to determine which driver(s) will manage the controller.

For reference, the following is the definition of **ConnectController()**:

```
typedef
EFI_STATUS
ConnectController (
  IN EFI_HANDLE                 ControllerHandle,
  IN EFI_HANDLE                *DriverImageHandle    OPTIONAL,
  IN EFI_DEVICE_PATH_PROTOCOL  *RemainingDevicePath  OPTIONAL,
  IN BOOLEAN                    Recursive
  );
```

The connection is a two-phase process, as follows:

- First phase: Construct an ordered list of driver handles.
- Second phase: Try to connect the drivers to a controller in priority order from first to last.

Table 2-8 lists the steps for phase one, which are known as the *driver connection precedence rules.* Much of this information is in section 5.3.1 of the *EFI 1.10 Specification* where the EFI Boot Service **ConnectController()** is discussed.

Table 2-8. Connecting Controllers: Driver Connection Precedence Rules

Step	Type of Override	Description
1	Context override	The parameter *DriverImageHandle* is an ordered list of image handles. The highest-priority driver handle is the first element of the list, and the lowest-priority driver handle is the last element of the list. The list is terminated with a **NULL** driver handle. This parameter is usually **NULL** and is typically only used to debug new drivers from the EFI Shell. These drivers are placed at the head of the ordered list of driver handles.
2	Platform driver override	If an **EFI_PLATFORM_DRIVER_OVERRIDE_PROTOCOL** instance is present in the system, then the **GetDriver()** service of this protocol is used to retrieve an ordered list of image handles for *ControllerHandle*. The first driver handle returned from **GetDriver()** has the highest precedence, and the last driver handle returned from **GetDriver()** has the lowest precedence. The ordered list is terminated when **GetDriver()** returns **EFI_NOT_FOUND**. It is legal for no driver handles to be returned by **GetDriver()**. Any driver handles obtained from the **EFI_PLATFORM_DRIVER_OVERRIDE_PROTOCOL** are appended to the end of the ordered list of driver handles from the context override (step #1 above).
3	Bus specific driver override	If there is an instance of the **EFI_BUS_SPECIFIC_DRIVER_OVERRIDE_PROTOCOL** attached to *ControllerHandle*, then the **GetDriver()** service of this protocol is used to retrieve an ordered list of driver handles for *ControllerHandle*. The first driver handle returned from **GetDriver()** has the highest precedence, and the last driver handle returned from **GetDriver()** has the lowest precedence. The ordered list is terminated when **GetDriver()** returns **EFI_NOT_FOUND**. It is legal for no driver handles to be returned from **GetDriver()**. Any driver handles obtained from the **EFI_BUS_SPECIFIC_DRIVER_OVERRIDE_PROTOCOL** are appended to the end of the ordered list of driver handles that was produced from the context override and platform driver override (steps 1 and 2). In practice, this precedent option allows the EFI drivers that are stored in an option ROM of a PCI adapter to manage the PCI adapter. Even if drivers with higher versions are available, this rule exists to make sure that if a particular driver on a card requires a hardware modification to go with it, the driver will not get replaced by later versions of the driver that may not work with the hardware modification.

continued

Table 2-8. Connecting Controllers: Driver Connection Precedence Rules (continued)

Step	Type of Override	Description
4	Driver binding search	The list of available driver handles can be found by using the boot service **LocateHandle()** with a *SearchType* of *ByProtocol* for the GUID of the **EFI_DRIVER_BINDING_PROTOCOL**. From this list, the driver handles found in steps 1, 2, and 3 above are removed. The remaining driver handles are sorted from highest to lowest based on the *Version* field of the **EFI_DRIVER_BINDING_PROTOCOL** instance that is associated with each driver handle. Any driver handles that are obtained from this search are appended to the end of the ordered list of driver handles started in steps 1, 2, and 3. In practice, this sorting means that a PCI adapter that does not have an EFI driver in its option ROM will be managed by the driver with the highest *Version* number.

Phase two of the connection process is to simply check each driver in the ordered list to see if it supports the controller. This check is done by calling the **Supported()** service of the driver's Driver Binding Protocol and passing in the *ControllerHandle* and the *RemainingDevicePath*. If successful, the **Start()** service of the driver's Driver Binding Protocol is called and passes in the *ControllerHandle* and *RemainingDevicePath*.. Each driver in the list is given an opportunity to connect, even if a prior driver connected successfully. However, if a driver with higher priority had already connected and opened the parent I/O protocol with exclusive access, the other drivers would not be able to connect.

One reason this type of connection process is used is because the order in which drivers are installed into the handle database is not deterministic. Drivers can be unloaded and reloaded later, which changes the order of the drivers in the handle database.

These precedent rules assume that the relevant drivers to be considered are loaded into memory. This case may not be true for all systems. Large systems, for example, may limit "bootable" devices to a subset of the total number of devices in the system.

The **ConnectController()** function can be called several times during the initialization of EFI. It is called before launching the boot manager to set the consoles.

2.9.2 Loading EFI Option ROM Drivers

The following is an interesting use case that tests these precedence rules. Assume that the following three identical adapters are in the system:

- Adapter A: EFI driver Version 0x10
- Adapter B: EFI driver Version 0x11
- Adapter C: No EFI driver

These three adapters have EFI drivers in the option ROM as defined below. When EFI connects, the drivers control the devices as follows:

- EFI driver Version 0x10 manages Adapter A.
- EFI driver Version 0x11 manages Adapter B and Adapter C.

If the user then goes into the EFI Shell and soft loads EFI driver version 0x12, nothing will change until the existing drivers are disconnected and a reconnect is performed. This reconnection can be done in a variety of ways, but the EFI Shell command **`reconnect -r`** is the easiest. Now the drivers control the devices as follows:

- EFI driver Version 0x10 manages Adapter A.
- EFI driver Version 0x11 manages Adapter B.
- EFI driver Version 0x12 manages Adapter C.

An OEM can override this logic by implementing the Platform Driver Override Protocol.

2.9.3 DisconnectController()

`DisconnectController()` performs the opposite of **`ConnectController()`**. It requests drivers that are managing a controller to release the controller.

2.10 Platform Initialization

Figure 2-7 shows the sequence of events that occur when an EFI-based system is booted. The following sections will describe each of these events in detail and how they relate to EFI drivers.

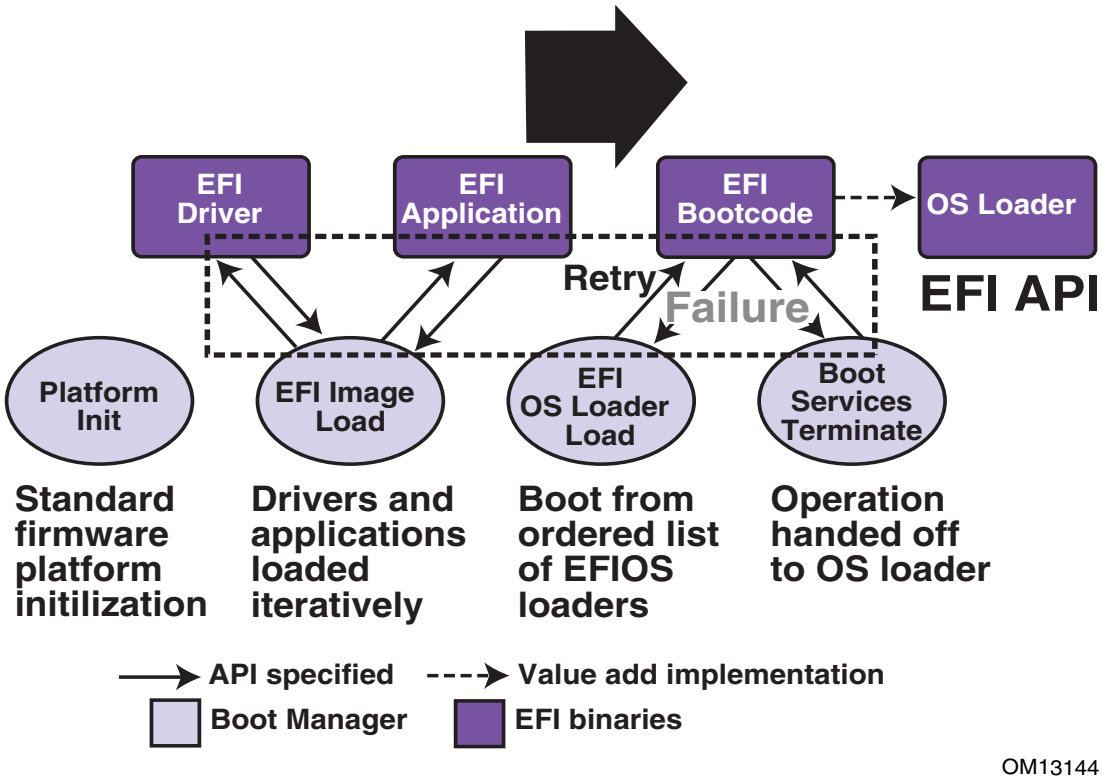

Figure 2-7. Booting Sequence

Figure 2-8 shows a possible system configuration. Each box represents a physical device (called a *controller*) in the system. Before the first EFI connection process is performed, none of these devices is registered in the handle database. The following sections describe the steps that the *EFI Sample Implementation* follows to initialize a platform and how drivers are executed, handles are created, and protocols are installed.

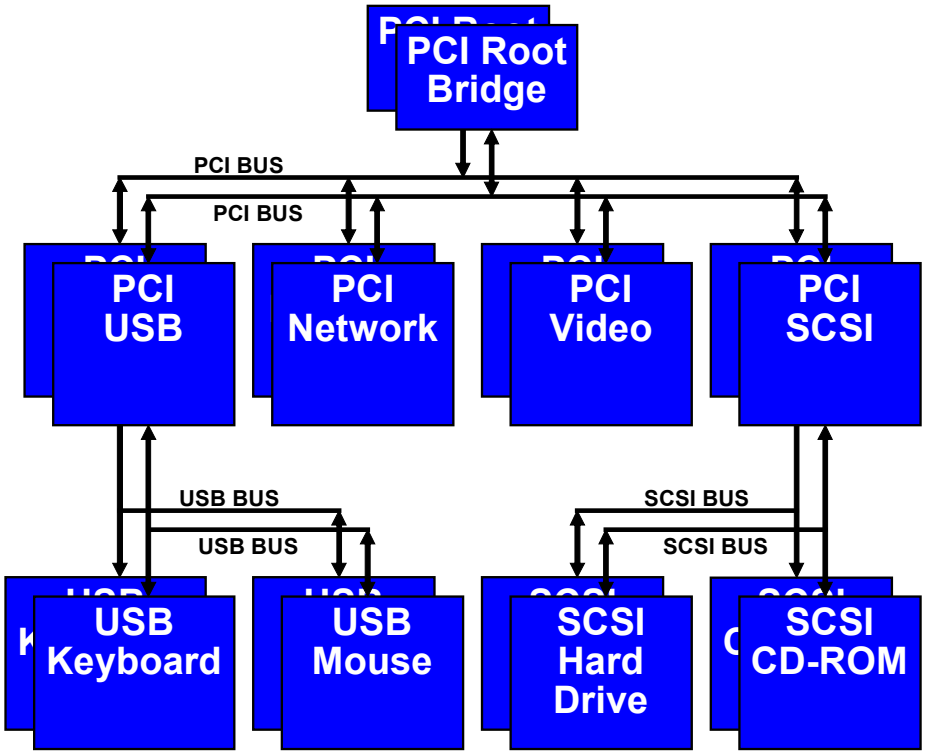

Figure 2-8. Sample System Configuration

At some point early in the boot process, EFI initialization will create handles and install the EBC Protocol and the Decompression Protocol in the handle database. These service protocols will be needed to run EFI drivers that may be compressed or compiled in an EBC format. For example, the handle database as viewed with the **dh** EFI Shell command might look like the following:

Handle Database

```
...
  6: Ebc
...
  9: Image(Decompress)
  A: Decompress
```

2.10.1 Connecting PCI Root Bridge(s)

During EFI initialization, the system will load a root bridge driver for the root device, typically the PCI root bridge driver, with **LoadImage()**. Like all drivers, as it is loaded, EFI will create a handle in the handle database and attach an instance of the **EFI_LOADED_IMAGE_PROTOCOL** with the unique image information for the PCI root bridge driver. Because this driver is the system root driver, it does not follow the EFI Driver Model. When it is loaded, it does not install the following:

- Driver Binding Protocol
- Component Name Protocol
- Driver Diagnostics Protocol
- Driver Configuration Protocol

Instead, it immediately uses its knowledge about the platform architecture to create handles for each PCI root bridge. This example shows only one PCI root bridge, but there can be many PCI root bridges in a system, especially in a data center server. It installs the **EFI_PCI_ROOT_BRIDGE_IO_PROTOCOL** and an **EFI_DEVICE_PATH_PROTOCOL** on each handle. By not installing the Driver Binding Protocol, the PCI root bridge prevents itself from being disconnected or reconnected later on. For example, the handle database as viewed with the **dh** EFI Shell command might look like the following after the PCI root bridge driver is loaded and executed. This example shows an image handle that is a single controller handle with a PCI Root Bridge I/O Protocol and the Device Path Protocol.

Handle Database

```
…
  B: Image(PcatPciRootBridge)
  C: DevIo PciRootBridgeIo DevPath (Acpi(HWP0002,0,PNP0A03))
```

NOTE

PNP0A03 may appear in either _HID or _CID of the PCI root bridge device path node. This example is one where it is not in _HID.

OS loaders need to access the boot devices to boot the OS. Such devices must have a Device Path Protocol, because the OS loader uses this protocol to determine the location of the boot device. Here at the root device, the Device Path Protocol contains a single ACPI node **Acpi(HID, UID)** or **Acpi(HID, UID, CID)**. This node points the OS to the place in the ACPI name space where the PCI root bridge is completely described. The **EFI_PCI_ROOT_BRIDGE_IO_PROTOCOL** provides PCI functions that are used by the PCI bus driver that is described in next section.

2.10.2 Connecting the PCI Bus

EFI initialization will continue by loading the PCI bus driver. As its entry point is executed, the PCI bus driver installs the Driver Binding Protocol and Component Name Protocol. For example, the handle database as viewed with the **dh** EFI Shell command might look like the following after the PCI bus driver is loaded and executed. It contains one new driver image handle with the Loaded Image Protocol, Driver Binding Protocol, and Component Name Protocol. Because this driver does follow the EFI Driver Model, no new controller handles are produced when the driver is loaded and executed. They will not be produced until the driver is connected.

Handle Database

```
…
14: Image(PciBus) DriverBinding ComponentName
```

At some point later in the initialization process, EFI will use **ConnectController()** to attempt to connect the PCI root bridge controller(s) (handle #C). The system has several priority rules for determining what driver to try first, but in this case it will search the handle database for driver handles (handles with the Driver Binding Protocol). The search will find handle #14 and call the driver's **Supported()** service, passing in controller handle #C. The PCI bus driver requires the Device Path Protocol and PCI Root Bridge I/O Protocol to be started, so the **Supported()** service will return **EFI_SUCESS** when those two protocols are found on a handle. After receiving **EFI_SUCCESS** from the **Supported()** service, **ConnectController()** will then call the **Start()** service with the same controller handle (#C).

Because the PCI bus driver is a bus driver, the **Start()** service will use the PCI Root Bridge I/O Protocol functions to search the PCI bus for all PCI devices. For each PCI device/function that the PCI bus driver finds, it will create a child handle and install an instance of the PCI I/O Protocol on the handle. The handle is registered in the handle database as a "child" of the PCI root bridge controller. The PCI driver will also copy the device path from the parent PCI root bridge device handle and append a new PCI device path node **Pci(Dev|Func)**. In cases where the PCI driver discovers a PCI-to-PCI bridge, the devices below the bridge will also be added as children to the bridge. In these cases, an extra PCI device path node will be added for each PCI-to-PCI bridge in between the PCI root bridge and the PCI device. For example, the handle database as viewed with the **dh** EFI Shell command might look like the following after the PCI bus driver is connected to the PCI root bridge. It shows the following:

- Nine PCI devices were discovered.
- The PCI device on handle #1B has an option ROM with an EFI driver.
- That EFI driver was loaded and executed and is shown as handle #1C.

Also notice that a single PCI card may have several EFI handles if they have multiple PCI functions.

Handle Database

```
…
16: PciIo DevPath (Acpi(HWP0002,0,PNP0A03)/Pci(1|0))
17: PciIo DevPath (Acpi(HWP0002,0,PNP0A03)/Pci(1|1))
18: PciIo DevPath (Acpi(HWP0002,0,PNP0A03)/Pci(2|0))
19: PciIo DevPath (Acpi(HWP0002,0,PNP0A03)/Pci(2|1))
1A: PciIo DevPath (Acpi(HWP0002,0,PNP0A03)/Pci(2|2))
1B: PciIo DevPath (Acpi(HWP0002,0,PNP0A03)/Pci(3|0))
1C: Image(Acpi(HWP0002,0,PNP0A03)/Pci(3|0)) DriverBinding
1D: PciIo DevPath (Acpi(HWP0002,0,PNP0A03)/Pci(4|0))
1E: PciIo DevPath (Acpi(HWP0002,100,PNP0A03)/Pci(1|0))
1F: PciIo DevPath (Acpi(HWP0002,100,PNP0A03)/Pci(1|1))
```

2.10.3 Connecting Consoles

At this point during the EFI initialization, EFI does not have a "console" device configured. This absence is because it is often a PCI device, and we have been waiting for the PCI bus driver to provide device handles for the console(s). Most EFI platforms will now follow a console connection strategy to connect the consoles in a consistent manner that is defined by the platform. This way the platform will be able to display messages to all of the selected consoles through the standard EFI mechanisms. Prior to this point, the platform messages, if any, were being displayed by platform-specific methods.

2.10.4 Console Drivers

Typical EFI consoles include the following:

- UGA
- USB keyboard
- USB mouse
- Serial ports

Some systems may have other devices. Table 2-9 on the next page shows an example of drivers that are installed in the handle database from the system ROM. These drivers are in the system ROM because they are core devices built into the motherboard. If the devices were PCI cards, the USB host controller driver and the UGA draw driver may have been found in the EFI option ROMs of those cards.

Table 2-9. EFI Console Drivers

Class of Driver	Type of Driver	Driver Name	Description and Example
USB Console	USB host controller driver	UsbOhci or UsbUhci	Consumes the PCI I/O Protocol and produces the USB Host Controller Protocol. `25: Image(UsbOhci) DriverBinding ComponentName`
	USB bus driver	UsbBus	Consumes the USB Host Controller Protocol and produces the USB I/O Protocol. `26: Image(UsbBus) DriverBinding ComponentName`
	USB keyboard driver	UsbKb	Consumes the USB I/O Protocol and produces the Simple Input Protocol. `27: Image(UsbKb) DriverBinding ComponentName`
	USB mouse driver	UsbMouse	Consumes the USB I/O Protocol and produces the Simple Pointer Protocol. `28: Image(UsbMouse) DriverBinding ComponentName`
UGA	UGA draw	CirrusLogic5430	Consumes the PCI I/O Protocol and produces the UGA Draw Protocol. `2E: Image(CirrusLogic5430) DriverBinding ComponentName`
	Graphics console driver	GraphicsConsole	Consumes the UGA Draw Protocol and produces the Simple Text Output Protocol. `2D: Image(GraphicsConsole) DriverBinding ComponentName`
Serial	Serial 16550 UART driver	Serial16550	Consumes the ISA I/O Protocol and produces the Serial I/O Protocol. `30: Image(Serial16550) DriverBinding ComponentName`
	Serial terminal driver	Terminal	Consumes the Serial I/O Protocol and produces the Simple Input and Simple Text Output Protocols. `31: Image(Terminal) DriverBinding ComponentName`

Generic Console	Platform console management driver	ConPlatform	This driver is a bit unique in that a single set of driver code adds two driver handles, one for the “Console Out” and another for the “Console In.”
			`32: Image(ConPlatform) DriverBinding ComponentName`
			`33: DriverBinding ComponentName`

	Console splitter driver	ConSplitter	This driver may not be present on all platforms. It will combine the various selected input and output devices for the following four basic EFI user devices: • ConIn • ConOut • ErrOut • PointerIn

continued

Table 2-9. EFI Console Drivers (continued)

Class of Driver	Type of Driver	Driver Name	Description and Example
Generic Console (continued)	Console splitter driver (continued)	ConSplitter (continued)	It also installs multiple driver handles for a single set of driver code. It installs driver handles to manage ConIn, ConOut, ErrOut, and PointerIn devices. The entry point of this driver creates virtual handles for ConIn, ConOut, and StdErr, respectively, that are called the following: • PrimaryConIn • PrimaryConOut • PrimaryStdErr The virtual handles will always exist even if no console exists in the system. These virtual handles are needed to support hot-plug devices such as USB. **34: Image(ConSplitter) DriverBinding ComponentName** **35: DriverBinding ComponentName** **36: DriverBinding ComponentName** **37: DriverBinding ComponentName** **38: Txtout PrimaryStdErr** **39: Txtin SimplePointer PrimaryConIn** **3A: Txtout PrimaryConOut UgaDraw**

2.10.5 Console Variables

After loading these drivers in the handle database, the platform can connect the console devices that the user has selected. The device paths for these consoles will be stored in the *ConIn*, *ConOut*, and *ErrOut* global EFI variables (see section 3.2 in the *EFI 1.10 Specification*). For the purpose of this example, the variables have the following device paths:

```
ErrOut = Acpi(HWP0002,0,PNP0A03)/Pci(1|1)/Uart(9600N81)/
         VenMsg(Vt100+);Acpi(HWP0002,0,PNP0A03)/Pci(4|0)

ConOut = Acpi(HWP0002,0,PNP0A03)/Pci(1|1)/Uart(9600N81)/
         VenMsg(Vt100+);Acpi(HWP0002,0,PNP0A03)/Pci(4|0)

ConIn  = Acpi(HWP0002,0,PNP0A03)/Pci(1|1)/Uart(9600N81)/
         VenMsg(Vt100+)
```

Note the following:

- The *ErrOut* and *ConOut* variables are compound device paths indicating that the EFI output is mirrored on two devices. This work is done by the console splitter driver when it is connected. The two devices are a serial terminal and a PCI video controller.
- The *ConIn* variable contains a device path to a serial terminal.
- The *ErrOut* variable is typically the same as the *ConOut* variable, but could be redirected to another device. It is important to check when developing EFI drivers because the **DEBUG()** output is typically directed to the *ErrOut* device and it may not always be the same device as the *ConOut* device. In this case, the two devices are a serial terminal and a PCI video controller.

2.10.6 ConIn

Now the system will connect the console devices using the device paths in the *ConIn*, *ConOut*, and *ErrOut* global EFI variables. The *ConIn* connection process will be discussed first.

```
ConIn  = Acpi(HWP0002,0,PNP0A03)/Pci(1|1)/Uart(9600 N81)/
         VenMsg(Vt100+)
```

The EFI connection process will search for the device in the handle database that has a device path the matches the following the closest.

```
Acpi(HWP0002,0,PNP0A03)/Pci(1|1)/Uart(9600 N81)/VenMsg(Vt100+)
```

It will find handle #17 as the closest match. The portion of the device path that did not match (**Uart(9600 N81)/VenMsg(Vt100+)**) is called the *remaining device path.*

```
17: PciIo DevPath (Acpi(HWP0002,0,PNP0A03)/Pci(1|1))
```

EFI will call **ConnectController()**, passing in handle #17 and the remaining device path. The connection code will construct a list of all the drivers in the system and call each driver, passing handle #17 and the remaining device path into the **Supported()** service. The only driver installed in the handle database that will return **EFI_SUCCESS** for this device handle is handle #30:

```
30: Image(Serial16550) DriverBinding ComponentName
```

Now that **ConnectController()** has found a driver that supports handle #17, it will pass device handle #17 and the remaining device path **Uart(9600 N81)/ VenMsg(Vt100+)** into the serial driver's **Start()** service. The serial driver will open the PCI I/O Protocol on handle #17 and create a new child handle. The following will be installed onto the new child handle:

- **EFI_SERIAL_IO_PROTOCOL** (defined in section 10.12 of the *EFI 1.10 Specification*)
- **EFI_DEVICE_PATH_PROTOCOL**

The device path for the child handle is generated by making a copy of the device path from the parent and appending the serial device path node **Uart(9600 N81)**. Handle #3B shown below is the new child handle.

```
3B: SerialIo DevPath (Acpi(HWP0002,0,PNP0A03)/Pci(1|1)/
      Uart(9600 N81))
```

That first call to **ConnectController()** has now been completed, but the Device Path Protocol on handle #3B does not completely match the *ConIn* device path, so the connection process will be repeated. This time the closest match for **Acpi(HWP0002,0)/Pci(1|1)/Uart(9600 N81)/VenMsg(Vt100+)** is the newly created device handle #3B. Now the remaining device path is **VenMsg(Vt100+)**. The search for a driver that supports handle #3B will find the terminal driver, returning **EFI_SUCCESS** from the **Supported()** service.

```
31: Image(Terminal) DriverBinding ComponentName
```

This driver's **Start()** service will open the **EFI_SERIAL IO_PROTOCOL**, create a new child handle, and install the following:

- **EFI_SIMPLE_INPUT_PROTOCOL**
- **EFI_SIMPLE_TEXT_OUTPUT_PROTOCOL**
- **EFI_DEVICE_PATH_PROTOCOL**

The console protocols are defined in sections 10.2 and 10.3 of the *EFI 1.10 Specification*. The device path will be generated by making a copy of the device path from the parent and appending the terminal device path node **VenMsg(Vt100+)**. VT100+ was chosen because that terminal type was specified in the remaining device path that was passed into the **Start()** service. Handle #3C show below is the new child handle.

```
3C: Txtin Txtout DevPath (Acpi(HWP0002,0,PNP0A03)/Pci(1|1)/
      Uart(9600 N81)/VenMsg(Vt100+))
```

The process still has not completely matched the *ConIn* device path, so the connection process will be repeated again. This time there is an exact match for **Acpi(HWP0002,0,PNP0A03)/Pci(1|1)/Uart(9600 N81)/VenMsg(Vt100+)** with the newly created child handle #3C. Searching for a driver that will support this controller will find two driver handles that return **EFI_SUCCESS** to the **Supported()** service. The two driver handles are from the platform console management driver:

```
32: Image(ConPlatform) DriverBinding ComponentName
33: DriverBinding ComponentName
```

Driver #32 will install a *ConOut* device GUID on the handle if the device path is listed in the *ConOut* global EFI variable. In our example, this case is true. Driver #32 will also install an *ErrOut* device GUID on the handle if the device path is listed in the *ErrOut* global EFI variable. This case is also true in our example. Therefore, handle #3C will have two new protocols on it: *ConOut* and *StdErr*.

```
3C: Txtin Txtout ConOut StdErr DevPath (..art(9600 N81)/
      VenMsg(Vt100+))
```

Driver #33 will install a *ConIn* device GUID on the handle if the device path is listed in the *ConIn* global EFI variable (which it will because we started this connection process that way), so handle #3C will have the *ConIn* protocol attached.

```
3C: Txtin Txtout ConIn ConOut StdErr DevPath (..art(9600
      N81)/VenMsg(Vt100+))
```

EFI uses these three protocols (*ConIn*, *ConOut*, and *ErrOut*) to mark devices in the system, which have been selected by the user as *ConIn*, *ConOut*, and *StdErr*. These protocols are actually just a GUID without any services or data.

There are three other driver handles that will return **EFI_SUCCESS** from the **Supported()** service. These driver handles are from the console splitter drivers for the *ConIn*, *ConOut*, and *ErrOut* devices in the system. There is a fourth console splitter driver handle (which is not used on this handle) for devices that support the Simple Pointer Protocol. The three driver handles are listed below:

```
34: Image(ConSplitter) DriverBinding ComponentName
36: DriverBinding ComponentName
37: DriverBinding ComponentName
```

Remember that when the console splitter driver was first loaded, it created three virtual handles. It marked these three handles with protocols *PrimaryStdErr*, *PrimaryConIn*, and *PrimaryConOut*. These protocols are only GUIDs. They do not contain any services or data.

```
38: Txtout PrimaryStdErr
39: Txtin SimplePointer PrimaryConIn
3A: Txtout PrimaryConOut UgaDraw
```

The console splitter driver's **Supported()** service for handle #34 examines the handle #3C for a *ConIn* Protocol. Having found it, it returns **EFI_SUCCESS**. The **Start()** service will then open the *ConIn* protocol on handle #3C such that the *PrimaryConIn* device handle #39 becomes a child controller and starts aggregating the **SIMPLE_INPUT_PROTOCOL** services.

The same thing happens for handle #36 with *ConIn*, except that the **SIMPLE_TEXT_OUTPUT_PROTOCOL** functionality on handle #3C is aggregated into the **SIMPLE_TEXT_OUTPUT_PROTOCOL** on the *PrimaryConOut* handle #3A.

Handle #37 with *StdErr* also does the same thing; the **SIMPLE_TEXT_OUTPUT_PROTOCOL** functionality on handle #3C is aggregated into the **SIMPLE_TEXT_OUTPUT_PROTOCOL** on the *PrimaryStdErr* handle #38.

The connection process has now been completed for *ConIn* because the device path that completely matched the *ConIn* device path and all the console-related services have been installed.

2.10.7 ConOut

As with *ConIn*, EFI will connect the *ConOut* devices using the device paths in the *ConOut* global EFI variable. If *ConIn* was not complicated enough, the *ConOut* global EFI device path in this example is a compound device path and indicates that the *ConOut* device is being mirrored with the console splitter driver to two separate devices.

```
ConOut = Acpi(HWP0002,0,PNP0A03)/Pci(1|1)/Uart(9600 N81)/
      VenMsg(Vt100+);Acpi(HWP0002,0,PNP0A03)/Pci(4|0)
```

The EFI connection process will search the handle database for a device path that matches the first device path in the *ConOut* variable:

```
Acpi(HWP0002,0,PNP0A03)/Pci(1|1)/Uart(9600 N81)/VenMsg(Vt100+)
```

Luckily, the device path already exists on handle #3C in its entirety thanks to the connection work done for *ConIn*.

```
3C: Txtin Txtout ConIn ConOut StdErr DevPath (..art(9600
      N81)/VenMsg(Vt100+))
```

EFI will perform a **ConnectController()** on handle #3C. Because this step was done previously with *ConIn*, there is nothing more to be done here.

The connection process has not yet been completed for *ConOut* because the device path is a compound device path and a second device needs to be connected:

```
Acpi(HWP0002,0,PNP0A03)/Pci(4|0)
```

The EFI connection process will search the handle database for a device path that matches **Acpi(HWP0002,0,PNP0A03)/Pci(4|0)**. The device path already exists in its entirety on handle #1C and was created by the PCI bus driver when it started and exposed the PCI devices.

```
1C: PciIo DevPath (Acpi(HWP0002,0,PNP0A03)/Pci(4|0))
```

EFI will perform a **ConnectController()** on handle #1C. Note that the device path is a complete match, so there is no remaining device path to pass in this time. **ConnectController()** constructs the prioritized list of drivers in the system and calls the **Supported()** service for each one, passing in the device handle #1C. The only driver that will return **EFI_SUCCESS** is the UGA driver.

```
2E: Image(CirrusLogic5430) DriverBinding ComponentName
```

ConnectController() will **Start()** this driver and it will consume the device's **EFI_PCI_IO_PROTOCOL** and install the **EFI_UGA_DRAW_PROTOCOL** onto the device handle #1C.

```
1C: PciIo UGADraw DevPath (Acpi(HWP0002,0,PNP0A03)/
      Pci(4|0))
```

ConnectController() will continue to process its list of drivers and will find that the "GraphicsConsole" driver's **Supported()** service will return **EFI_SUCCESS**.

```
2D: Image(GraphicsConsole) DriverBinding ComponentName
```

Next the graphics console driver's **Start()** service will consume the **EFI_UGA_DRAW_PROTOCOL** and produce the **EFI_SIMPLE_TEXT_OUTPUT_PROTOCOL** on the same device handle #1C.

```
1C: Txtout PciIo UgaDraw DevPath (Acpi(HWP0002,0,PNP0A03)/
      Pci(4|0))
```

ConnectController() will continue to process its list of drivers. This time, searching for a driver that will support this controller will find two driver handles that return **EFI_SUCCESS** to the **Supported()** service. These two driver handles are from the platform console management driver:

```
32: Image(ConPlatform) DriverBinding ComponentName
```

Driver handle #32 will install a *ConOut* device GUID on the handle if the device path is listed in the *ConOut* global EFI variable. In this example, this case is true. Driver #32 will also install an *ErrOut* device GUID on the handle if the device path is listed in the *ErrOut* global EFI variable. This case is also true in the example. Therefore, handle #1C has two new protocols on it: *ConOut* and *ErrOut*.

```
1C: Txtout PciIo ConOut StdErr DevPath
      Acpi(HWP0002,0,PNP0A03)/Pci(4|0))
```

These two protocols (*ConOut* and *ErrOut*) are used to mark devices in the system that have been selected by the user as *ConOut* and *StdErr*. These protocols are actually just a GUID without any functions or data.

There are two other driver handles that will return **EFI_SUCCESS** to the **Supported()** service. These driver handles are from the console splitter driver for the *ConOut* and *ErrOut* devices in the system.

```
36: DriverBinding ComponentName
37: DriverBinding ComponentName
```

Remember that when the console splitter driver was first loaded, it created three virtual handles. It marked these three handles with protocols *PrimaryStdErr*, *PrimaryConIn*, and *PrimaryConOut*. These protocols are only GUIDs. They do not contain any services or data.

```
38: Txtout PrimaryStdErr
39: Txtin SimplePointer PrimaryConIn
3A: Txtout PrimaryConOut UgaDraw
```

The console splitter driver's **Supported()** service for driver handle #36 examines the handle #1C for a *ConOut* Protocol. Having found it, it returns **EFI_SUCCESS**. The **Start()** service will then open the *ConOut* protocol on device handle #1C such that the *PrimaryConOut* device handle #3A becomes a child controller and starts aggregating the **SIMPLE_TEXT_OUTPUT_PROTOCOL** services.

The same thing happens for driver handle #37 with *StdErr*; the **SIMPLE_TEXT_OUTPUT_PROTOCOL** functionality on device handle #1C is aggregated into the **SIMPLE_TEXT_OUTPUT_PROTOCOL** on the *PrimaryStdErr* device handle #38.

2.10.8 ErrOut

In this example, *ErrOut* is the same as *ConOut*. So the connection process for *ConOut* is executed one more time.

```
ErrOut = Acpi(HWP0002,0,PNP0A03)/Pci(1|1)/Uart(9600 N81)/
      VenMsg(Vt100+);Acpi(HWP0002,0,PNP0A03)/Pci(4|0)
```

2.10.9 Loading Other Core Drivers

After connecting the consoles, most platforms will typically load some additional core drivers, but this step is dependent on platform policy. For example, the *EFI Sample Implementation* loads the following drivers after consoles are connected:

- Disk media support drivers (DiskIo, Partition, Fat)
- Platform media drivers (for example, platform-specific bus drivers for SCSI and IDE; for this example, it will load "XyzScsi")
- LAN support
- Core LAN drivers (XyzUndi)
- LAN support drivers (BiosSnp16, Snp3264, PxeBc, PxeDhcp4, BIS)
- FPSWA (Itanium architecture only)

Remember that these drivers are only loaded. They are not connected to controllers until **ConnectController()** is called.

2.10.10 Boot Manager Connect ALL

On some platforms, the EFI boot manager may connect all drivers to all devices at this point in the platform initialization sequence. However, platform firmware can choose to connect the minimum number of drivers and devices that is required to establish consoles and gain access to the boot device. Performing the minimum amount of work is recommended to enable shorter boot times.

If the platform firmware chooses to go into a "platform configuration" mode, then all the drivers should be connected to all devices. The following algorithm is used to connect all drivers to all devices.

A search is made of the handle database for all root controller handles. These handles do not have a Driver Binding Protocol or the Loaded Image Protocol. They will have a Device Path Protocol, and they will not have any parent controllers. **ConnectController()** is called with the *Recursive* flag set to **TRUE** and a *RemainingDevicePath* of **NULL** for each of the root controllers. These settings cause all children to be produced by all bus drivers. As each bus is

expanded, if the bus supports storage devices for EFI drivers, then additional EFI drivers will be loaded from those storage devices (for example, option ROMs on PCI adapters). This process is recursive. Each time a child handle is created, **ConnectController()** is called again on that child handle, so all of those handle's children will be produced. When the process is complete, then the entire tree of boot devices in the system hierarchy will be present in the handle database.

2.10.11 Boot Manager Driver Option List

The EFI boot manager loads the drivers that are specified by the *DriverOrder* and *Driver####* environment variables. These environment variables are discussed in sections 3.1 and 3.2 of the *EFI 1.10 Specification*. The EFI shell command **bcfg** can be used to add, remove, and display drivers to this list.

Before the EFI boot manager loads each driver, it will use the device path stored in the *Driver####* variable to connect the controllers and drivers that are required to access the driver option. This process is exactly the same as the process used for the console variables *ErrOut*, *ConOut*, and *ConIn*.

If any driver in the *DriverOrder* list has a load attribute of **LOAD_OPTION_FORCE_RECONNECT**, then the EFI boot manager will use the **DisconnectController()** and **ConnectController()** boot services to disconnect and reconnect all the drivers in the platform. This load attribute allows the newly loaded drivers to be considered in the driver connection process.

For example, if no driver in the *DriverOrder* list has the **LOAD_OPTION_FORCE_RECONNECT** load attribute, then it would be possible for a built-in system driver with a low version number to manage a device. Then, after loading a newer driver with a higher version number from the *DriverOrder* list, the driver with the lower version number will still manage the same device. However, if the newer driver in the *DriverOrder* list has a load attribute of **LOAD_OPTION_FORCE_RECONNECT**, then the EFI boot manager will disconnect and reconnect all the controllers, so the driver with the highest version number will manage the same device that the lower versioned driver was managing. Drivers that are added to the *DriverOrder* list should not set the **LOAD_OPTION_FORCE_RECONNECT** attribute unless they have to because the disconnect and reconnect process will increase the boot time.

2.10.12 Boot Manager BootNext

After connecting any drivers in the *DriverOrder* list, the EFI boot manager will attempt to boot the boot option that is specified by the *BootNext* environment variable. This environment variable is discussed in sections 3.1 and 3.2 of the *EFI 1.10 Specification*. This variable typically is not set, but if it is, EFI will delete the variable and then attempt to load the boot option that is described in the *Boot####* variable pointed to by *BootNext*.

Before the EFI boot manager boots the boot option, it will use the device path stored in the *Boot####* variable to connect the controllers and drivers that are required to access the boot option. This process is exactly the same as the process that is used for the console variables *ErrOut*, *ConOut*, and *ConIn*.

2.10.13 Boot Manager Boot Option List

The EFI boot manager will display the boot option menu and if the auto-boot `TimeOut` environment variable has been set, then the first boot option will be loaded when the timer expires. The boot options can be enumerated by the EFI boot manager by reading the `BootOrder` and `Boot###` environment variables. These environment variables are discussed in sections 3.1 and 3.2 of the *EFI 1.10 Specification*. A boot option is typically an OS loader that never returns to EFI, but boot options can also be EFI applications like diagnostic utilities or the EFI Shell. If a boot option does return to the EFI boot manager, and the return status is not **`EFI_SUCCESS`**, then the EFI boot manager the next boot option will be processed. This process is repeated until an OS is booted, **`EFI_SUCESS`** is returned by a boot option or the list of boot options is exhausted. Once the boot process has halted, the EFI boot manager will typically provide a user interface that allows the user to manually boot an OS or manage the platform.

The EFI boot manager will use the device path in each boot option to ensure that the device required to access the boot option has been added to the EFI handle database. This process is exactly the same as the process used for the console variables `ErrOut`, `ConOut`, and `ConIn`.

3
Coding Conventions

This section describes the coding conventions that are used in the *EFI Sample Implementation*. Most of the code that is distributed in the *EFI Sample Implementation* is implemented in the C programming language. One goal of the EFI Driver Model is to allow portable EFI drivers to be designed such that they can be easily compiled for various processor architectures. The processor architectures supported today include the following:

- IA-32
- Itanium architecture
- EFI Byte Code (EBC)

The EFI Byte Code virtual machine architecture was optimized with the C programming language in mind, so if an EFI driver writer wishes to design for portability, then the EFI driver must be implemented in the C programming language with no use of assembly language. As a result, the coding convention presented here concentrates on conventions for the C programming language. This coding convention is also used in the code examples presented throughout this document.

EFI drivers are not required to use these coding conventions. There is no data that suggests that one code convention is better than another. However, there is a large body of data that suggests that the use of a consistent coding convention improves the efficiency of software development. There are many advantages to using a common coding convention, including the following:

- Source code is easier to read.
- Source code is easier to maintain.
- It improves collaboration between developers.
- EFI driver are easier to debug.

Recommendations are presented for the naming of files, functions, macros, data types, variables, and constants. A consistent naming convention helps improve the readability of the code and it tends to make the code somewhat self documenting, which may help reduce the required comment overhead. In addition, a consistent function naming convention improves the readability of call stacks during debug. Conventions for the use of tabs, spaces, indentation, and comments are also presented along with templates for the various expressions and constructs that are available in the C programming language. Finally, the special considerations for implementing new EFI protocols, EFI GUIDs, and EFI drivers are presented.

3.1 Indentation and Line Length

Tabs are not used in any of the source files. Instead, spaces are used for all indentation. All code blocks are indented in increments of 2 spaces. This 2 space indentation is smaller than some other conventions that use 4 spaces, but this reduced indentation allows more code to be viewed on a single screen even when fairly deep nesting is used.

In general, the length of a line of code does not exceed 80 characters. This length is only a guideline, however. The main reason for defining a guideline on line length is so code can be easily viewed in a full-screen editor and formatted correctly when it is sent to a printer.

3.2 Comments

3.2.1 File Headers Comments

Every **.c** and **.h** file has a file header comment block at the very top of the file. The file header comment block has the form shown in Example 3-1.

```
/*++

Copyright (c) 2003 Xyz Corporation

Module Name:

  <<file name, e.g. Foo.c>>

Abstract:

  <<description of file contents>>

--*/
```

Example 3-1. File Header Comment Block for .C and .H Files

Files with an extension of **.inf** are essentially a shorthand **makefile** that the build tools in the *EFI Sample Implementation* use to determine the source files and libraries that are required to build an EFI driver. Every **.inf** file has a file header comment block at the very top of the file as shown in Example 3-2.

```
#
# Copyright (c) 2003 Xyz Corporation
#
# Module Name:
#
#   <<file name, e.g. Make.inf>>
#
# Abstract:
#
#   <<description of the file contents, e.g. Makefile for Edk\Drivers\DiskIo>>
#
```

Example 3-2. File Header Comment Block for .INF Files

3.2.2 Function Header Comments

Each function has a comment block between the parameter declaration section of the function definition and the opening brace of the function body. The only exceptions are simple or small private functions. The function header comment block takes the form as shown in Example 3-3.

```
EFI_STATUS
Foo (
  IN UINTN  Arg1,
  IN UINTN  Arg2
  )
/*++

Routine Description:

  <<description>>

Arguments:

  <<argument names and purposes>>

Returns:

  <<description of possible return values>>

--*/
{
}
```

Example 3-3. Function Header Comment Block

3.2.3 Internal Comments

Internal code comments use C++ style (**//**) comment lines. A block of comment lines that contain text are preceded and followed by a blank comment line. A blank line may optionally follow a comment block. The blank line generally indicates that the comment is for a large block of code. If a comment block is not followed by a blank line, then the comment only applies to the next few lines of code. The comments are indented with the code.

Comments are not placed on the same line as source code, but comments may be on the same line as data structures declarations and data structure initializations. Commenting is done as efficiently as possible because too many comments are as bad as too few. Comments are added to explain why things are done and the big picture of how a module works. Example 3-4 shows examples of internal comments.

```
  //
  // This is an example comment for the large block of code below
  //

  if (Test) {
    //
    // This is an example comment for the next code line
    //
    return Test;
  }
```

Example 3-4. Internal Comments

3.2.4 What Is Commented

The following things are commented in source code:

- Complicated, tricky, or sensitive pieces of code. Making the code cleaner is often better than adding more comments.
- Higher-level concepts in the code. Focus on the why and not the how.
- Data structures and `#define` statements in include files.
- The version number of any industry standard that is referenced by the code.

3.2.5 What Is Not Commmented

The following things are not commented in source code:

- In general, a single line of code does not need a comment.
- Where possible, the code is self documenting. Do not repeat the code or explain it in a comment. Instead, comments clarify the intent of the code or explain higher-level concepts.
- Do not place markers in the code, such as a developer's name or unusual patterns. They may have meaning to the developer but do not to other developers or projects.
- Code is not removed or disabled by using comments. Instead, a source control system is used to maintain different source code file revisions.

3.3 General Naming Conventions

Good naming practice is used when declaring variable names. Studies show that programs with names that average 10 to 16 characters in length are the easiest to debug (*Code Complete*). Programs with names averaging 8 to 20 characters are almost as easy to debug. This name length is simply a guideline. The most important thing is that the name conveys a clear meaning of what it represents.

Names are formed with a mix of upper and lower case text. Each word or concept starts with a capital letter and is followed by lower case letters for the rest of the word. Only the first letter of an acronym is capitalized. For example:

```
ThisIsAnExampleOfWhatToDoForPci
```

Commonly used terms are not be overloaded when possible. For example, EFI has an event model, so creating a new abstraction called an "event" would have caused confusion. The goal is to use names such that other developers with very little context can understand what each name represents.

3.3.1 Abbreviations

Table 3-1 describes a common set of abbreviations. Abbreviations do not have to be used, but if abbreviations are used, then the common one is selected. If new abbreviations are required, then making the abbreviation easy to remember is more important to the reader of the code than the writer of the code. New abbreviations are declared in a comment block at the top of the file. The translation table contains the new abbreviations and definitions. An attempt is made to not define new abbreviations to replace an existing abbreviation. For example, *No* is not redefined to mean *Number*, because *Num* is already defined as the supported abbreviation.

Table 3-1. Common Abbreviations

Abbreviation	Description
Ptr	Pointer
Str	Unicode string
Ascii	ASCII string
Len	Length
Arg	Argument
Max	Maximum
Min	Minimum
Char	Character
Num	Number
Temp	Temporary
Src	Source
Dst	Destination
BS	EFI Boot Services Table
RT	EFI Runtime Table
ST	EFI System Table
Tpl	EFI Task Priority Level

3.3.2 Acronyms

The use of acronyms is limited. The idea is to code for a developer who will have to read and maintain the code. Making up a new vernacular to describe a module can lead to considerable confusion. If new acronyms are created, then they are fully defined in the documentation and each module that uses the acronym contains a comment with a translation table in the file header.

The use of acronyms for industry standards is acceptable. Acronyms such as Efi, Pci, Acpi, Smbios, and Isa (capitalized to the variable naming convention) are used without defining their meaning. If an industry standard acronym is referenced, then the file header defines which revision of the specification is being assumed. For example, a PCI resource manager would state that it was coded to follow the PCI 2.2 specification.

3.4 Directory and File Names

Directory names and file names follow the general naming conventions that were presented in the previous section. There are no artificial limits on the length of a directory name or a source file name. None of the build tools in the *EFI Sample Implementation* require the file names to follow the 8.3 naming convention. New code is added below the **Edk** directory in the *EFI 1.10 Sample Implementation*. Example 3-5 contains some example directory names and file names from the area of the source tree where new drivers, protocols, GUIDs, and libraries would be placed during development. This example shows the source files for the following:

- The PCI video driver that produces the **EFI_UGA_DRAW_PROTOCOL**
- The declaration of the **EFI_BLOCK_IO_PROTOCOL**
- **EFI_GLOBAL_VARABLE_GUID**, which is used to access EFI variables
- A few source files from the EFI Driver Library

It is important for these directory and file name conventions to be followed when directories and files are created and when directory or file names are referenced from source files. These conventions will provide compatibility with case-sensitive file systems.

```
Efi1.1\
  Edk\
    Drivers\
      CirrusLogic5430\
        CirrusLogic5430.c
        CirrusLogic5430.h
        CirrusLogic5430UgaDraw.c
        ComponentName.c
        Make.inf
    Protocol\
      BlockIo\
        BlockIo.c
        BlockIo.h
    Guid\
      GlobalVariable\
        GlobalVariable.c
        GlobalVariable.h
    Lib\
      EfiDriverLib\
        EfiDriverLib.c
        Make.inf
```

Example 3-5. Directory and File Names

3.5 Function Names and Variable Names

Data variables and function names follow the general naming conventions. Each word or concept starts with a capital letter and is followed by lowercase letters. Macros, **#define**, and **typedef** declarations are all capital letters. Each word is separated by the underscore '**_**' character. The use of '**_t**' or '**_T**' is discouraged. Example 3-6 shows some example function and variable names.

```
#define THIS_IS_AN_EXAMPLE_OF_WHAT_TO_DO_FOR_PCI 1

#define MY_MACRO(a) ((a) == 1)

typedef struct {
  UINT32  Age;
  CHAR16  Name[32];
} MY_STRUCTURE;

EFI_STATUS
GetStructure (
  MY_STRUCTURE  *MyStructure
)
{
  EFI_STATUS  Status;
  . . .
}
```

Example 3-6. Function and Variable Names

3.5.1 Hungarian Prefixes

The Hungarian naming convention is generally discouraged in the coding conventions described here except for the global variable. The Hungarian naming convention is a set of detailed guidelines for naming variables and routines. The convention is widely used with the C programming language. The term "Hungarian" refers both to the fact that the names that follow the convention look like words from a foreign language and that the creator of the convention, Charles Simoyi, is originally from Hungary. The following is an example of a variable named with the "Hungarian" conventions:

pachInsert - A pointer to an array of characters to insert.

Hungarian conventions are not recommended for the following reasons:

- The abstraction of abstract data types is not covered. Instead, base types based on C programming language type–like integers and long integers are abstracted. As a result, the names are focused on data types instead of the object-oriented abstractions that they represent. This focus is of little value and forces manual type checking that can just as easily be accomplished by using a compiler with strong type checking or other source code analysis tools.
- Hungarian combines data meaning with data representation. If a change is made to a data type, then all the variables of that data type have to be renamed. In addition, there is no mechanism to ensure that the names are accurate.
- Studies have shown that Hungarian notation tends to encourage lazy variable names (*Code Complete*). It is common for developers to focus on the Hungarian prefix without putting effort into a good descriptive name.

3.5.2 Global Variables and Module Variables

The only exception to the Hungarian prefix model is the prefixing of a global variable with a 'g' and a module variable with an 'm.' For example:

```
gThisIsAGlobalVariableName
mThisIsAModuleVariableName
```

The use of global data is discouraged by this coding convention. The use of module variables is appropriate for solving specific programming issues. A *module* is defined to be a set of data and routines that act on that data. Thus, an EFI driver that produces a protocol can be thought of as a module. A complex protocol may be built out of several smaller modules. Any variable with scope outside of a single routine is prefixed by an 'm' or a 'g.'

Some module variables and global variables will require a locking mechanism to guarantee that the variable is being accessed by only one agent at a time. The locking mechanism that is available to EFI drivers is discussed in detail in chapter 4 of this document.

The difference between a module variable and a global variable is not always obvious. A module variable is accessed only by a small set of routines, typically within a single source file, that have strict rules for accessing the module variable. A global variable is accessed throughout the program, typically from multiple source files. Accessing a module variable would not be common over the life cycle of a program. New code, bug fixes, or new features would access the current routines that are used to abstract the module variable. The module variable would be easy to change, because it is accessed only from a small number of routines within a single source file. On the other hand, a global variable is accessed throughout the EFI driver and as the EFI driver evolves, more code will tend to access the global variable. Global variables and module variables are valid if they are being used to implement good information that is hiding in the design.

3.6 Macro Names

Macros use a different naming convention than functions. The main reason is the difference in the order of precedence that can occur between poorly constructed macros and functions. An overuse of parentheses is strongly encouraged, because it is very difficult to debug a precedence order bug in a macro. The following are examples of macro construction:

```
#define BAD_MACRO(a, b) a*b
#define GOOD_MACRO(a, b) ((a)*(b))
```

The following examples show the difference between **`BAD_MACRO()`** and **`GOOD_MACRO()`**:

- **`BAD_MACRO (10, 2)`** and **`GOOD_MACRO (10, 2)`** both evaluate to 20.
- **`BAD_MACRO (7+3, 2)`** evaluates to 7 + 3 * 2 = 7 + (3 * 2) = 7 + 6 = 13.
- **`GOOD_MACRO (7+3, 2)`** evaluates to (7+3) * (2) = 10 * 2 = 20.

3.6.1 Macros as Functions

It is recommended that all macros be constructed using the macro rules. Even with these recommendations, it is possible for a macro and an equivalent function to behave differently. All macros are designed so that they can be easily converted to an equivalent function implementation. Example 3-7 shows both good and bad examples of macros and their equivalent function implementations.

```
#define BAD_MACRO(Value)  ((Value) == 1) || ((Value) == 2)

#define GOOD_MACRO(Value)  ((Value) == 1)

BOOLEAN
BadMacro (
  UINTN Value
)
{
  return ((Value == 1) || (Value == 2));
}

BOOLEAN
GoodMacro (
  UINTN  Value
)
{
  return (Value == 1);
}
```

Example 3-7. Macros as Functions

When **GOOD_MACRO(Value)** and **GoodMacro(Value)** are called, the results are identical.

When **BAD_MACRO(Value)** and **BadMacro(Value)** are called, the results are identical.

When **GOOD_MACRO(Value++)** and **GoodMacro(Value++)** are called, the results are identical.

When **BAD_MACRO(Value++)** and **BadMacro(Value++)** are called, the results are not identical when *Value* is zero. Also, when **BAD_MACRO(Value++)** is called, *Value* is incremented twice, and when **BadMacro(Value++)** is called, *Value* is incremented only once.

3.7 Data Types

The EFI coding convention defines a set of common data types that are used to ensure portability between different compilers and processor architectures. Any abstract type that is defined is constructed from other abstract types or from common EFI data types. The types **int**, **unsigned**, **char**, **void**, **static**, and **long** are not used in this coding convention.

There is one exception to the preceding rules. The use of **static** is used for data but not for functions. The **STATIC** keyword is used to disable the static nature of function names to make their scope global to aid in debugging. This conversion does not influence the behavior of the program. Converting a **static** data type to nonstatic types can impact the way a program functions and thus the use of **static** for data is allowed.

Table 3-2 contains common data types that are referenced in the interface definitions defined in the *EFI 1.10 Specification*. Unless otherwise specified, all data types are naturally aligned. Structures are aligned on boundaries that are equal to the largest internal datum of the structure, and internal data is implicitly padded to achieve natural alignment.

The mappings from these compiler-dependent types to the EFI base types are handled in a single include file called **EfiBind.h**. If a new compiler is used, then **EfiBind.h** may need to be updated, but once the mapping between the new compiler's types and the EFI base types is made, all existing EFI source code should compile correctly.

Table 3-2. Common EFI Data Types

Mnemonic	Description
BOOLEAN	Logical Boolean. 1-byte value containing a 0 for **FALSE** or a 1 for **TRUE**. Other values are undefined.
INTN	Signed value of native width (4 bytes on IA-32, 8 bytes on Itanium architecture operations).
UINTN	Unsigned value of native width (4 bytes on IA-32, 8 bytes on Itanium architecture operations).
INT8	1-byte signed value.
UINT8	1-byte unsigned value.
INT16	2-byte signed value.
UINT16	2-byte unsigned value.
INT32	4-byte signed value.
UINT32	4-byte unsigned value.
INT64	8-byte signed value.
UINT64	8-byte unsigned value.
CHAR8	1-byte character.
CHAR16	2-byte character. Unless otherwise specified, all strings are stored in the UTF-16 encoding format as defined by Unicode 2.1 and ISO/IEC 10646 standards.
VOID	Undeclared type.
EFI_GUID	128-bit buffer containing a unique identifier value. Unless otherwise specified, aligned on a 64-bit boundary.
EFI_STATUS	Status code. Type INTN.
EFI_HANDLE	A collection of related interfaces. Type VOID *.
EFI_EVENT	Handle to an event structure. Type VOID *.

continued

Table 3-2. Common EFI Data Types (continued)

Mnemonic	Description
EFI_LBA	Logical block address. Type UINT64.
EFI_TPL	Task priority level. Type UINTN.
EFI_MAC_ADDRESS	32-byte buffer containing a network Media Access Controller address.
EFI_IPv4_ADDRESS	4-byte buffer. An IPv4 internet protocol address.
EFI_IPv6_ADDRESS	16-byte buffer. An IPv6 internet protocol address.
EFI_IP_ADDRESS	16-byte buffer aligned on a 4-byte boundary. An IPv4 or IPv6 internet protocol address.
<Enumerated Type>	Element of an enumeration. Type INTN.

Table 3-3 defines modifiers that are used in function and data declarations. The **IN**, **OUT**, and **OPTIONAL** modifiers are used only to qualify arguments to functions and therefore do not appear in data type declarations. The **IN** and **OUT** modifiers are placed at the beginning of the line before the data type of each function argument, and the **OPTIONAL** modifier is placed at the end of the line after the name of the function argument. The **STATIC** modifier is used to modify the scope of a function and can be overloaded to support debugging. The **EFIAPI** modifier is used to ensure that the correct calling convention that is defined in the *EFI Specification* is used between different modules that are not linked together.

Table 3-3. Modifiers for Common EFI Data Types

Mnemonic	Description
IN	Datum is passed to the function. Placed at the beginning of a source line before the data type of the function argument.
OUT	Datum is returned from the function. Placed at the beginning of a source line before the data type of the function argument.
OPTIONAL	Datum that is passed to the function is optional, and a **NULL** may be passed if the value is not supplied. Placed at the end of a source line after the name of the function argument.
STATIC	The function has local scope. This modifier replaces the standard C static key word, so it can be overloaded for debugging.
VOLATILE	Declares a variable to be volatile and thus exempt from optimization to remove redundant or unneeded accesses. Any variable that represents a hardware device is declared as **VOLATILE**.
CONST	Declares a variable to be of type **const**. This modifier is a hint to the compiler to enable optimization and stronger type checking during compile time.
EFIAPI	Defines the calling convention for EFI interfaces. All EFI intrinsic services and any member function of a protocol use this modifier on the function definition.

3.7.1 Enumerations

Example 3-8 is an example of a construction for an enumerated type. The names of the elements in the enumerated type follow the same naming convention as variables and functions. The **enum** is declared as a **typedef** with the name of the **typedef** following the standard all-capitalization rules. It is recommended that the last element of the **enum** be a member element that represents the maximum legal value for the enumeration. This representation allows for bounds checking on an **enum** to support debugging and sanity checking the value assigned to an **enum**. It is also recommended that the **enum** members be named carefully so that their names do not collide with other variable or function names.

```
typedef enum {
  EnumMemberOne,
  EnumMemberTwo,
  EnumMemberMax
} ENUMERATED_TYPE;
```

Example 3-8. Enumerated Types

3.7.2 Data Structures and Unions

Example 3-9 is an example of a construction of a data structure and a union. The names of the field in the data structure and the union follow the same naming conventions as variable and functions. A **struct** or **union** is declared as a **typedef** with the name of the **typedef** following the standard all-capitalization rules.

```
typedef struct {
  UINT32  FieldOne;
  UINT32  FieldTwo;
  UINT32  FieldThree;
} MY_STRUCTURE;

typedef union {
  UINT16  Integer;
  CHAR16  Character;
} MY_UNION;
```

Example 3-9. Data Structure and Union Types

3.8 Constants

Table 3-4 lists the constants that are available to all EFI drivers.

Table 3-4. EFI Constants

Mnemonic	Description
TRUE	One
FALSE	Zero
NULL	**VOID** pointer to zero.

3.9 Include Files

Include files contain a **#ifndef** and a **#define** statement at the start of the include file and a **#endif** statement as the last line of the file. The **#ifndef** and **#define** statements contain an all-uppercase version of the include file name, prefixed and postfixed by the '**_**' character, which prevents duplicate definitions if the same include file is included by a different module. All C include files use the **.h** file extension. The example in Example 3-10 shows the include file **SerialDriver.h**. It contains **#ifndef** and **#define** statements of **_EFI_SERIAL_DRIVER_H_**.

```
/*++

Copyright (c) 2003 Xyz Corporation

Module Name:

  SerialDriver.h

Abstract:

  Serial driver ….

--*/

#ifndef _EFI_SERIAL_DRIVER_H_
#define _EFI_SERIAL_DRIVER_H_

//
// Statements that include other header files
//

//
// Simple defines of such items as status codes and macros
//

//
// Type declarations
//

//
// Function prototype declarations
//

#endif
```

Example 3-10. Include File

The above comments suggest an order of declarations in an include file. Include files contain public declarations or private declarations, but not both. Include files do not contain code or declare storage for variables. Examples of public include files would be protocol definitions or industry standard specifications (such as EFI, ACPI, and SMBIOS). Private include files would contain static functions and internal data structure definitions.

3.10 Spaces in C Code

Spacing guidelines are very important and greatly improve the readability of source code. The program shown in Example 3-11 is an example that uses poor spacing guidelines and is very difficult to follow (Dishonorable mention, Obfuscated C Code Contest, 1984. Author requested anonymity).

```
int i;main(){for(;i["]<i;++i){--i;}"];read('-'-'-',i+++"hello,world!\n",'/'/'/'
));}read(j,i,p){write(j/p+p,i---j,i/i);}
```

Example 3-11. Poor Spacing

3.10.1 Vertical Spacing

Blank lines are used to make code more readable and to group logically related sections together. The convention of one statement per line is followed. This convention is not followed in **for** loops, where the initial, conditional, and loop statements reside on a single source line. The **if** statement is not an exception to this guideline. The conditional expression and body go on separate lines. The open brace (**{**) is placed at the end of the first line of the construct. The close brace (**}**) and **case** labels are placed on their own line. The close brace does share a line with the **else** and **else if** constructs.

3.10.2 Horizontal Spacing

Spaces are placed around assignment operators, binary operators, after commas, and before an open parenthesis. Spaces are not placed around structure members, pointer operators, or before open brackets of array subscripts. It is also better to use extra parentheses in more complex expressions rather than to depend on in-depth knowledge of the precedence rules of C. In addition, continuation lines are formatted so that they line up with the portion of the preceding line that they continue. Example 3-12 contains several practical examples of these spacing conventions.

```
Status = TestString (String, Index + 3, &Value);
MinResult = MIN_EXAMPLE (Start, End);
Data.Index
Pointer->Index = *Ptr;
Array[(Max + Min) / 2]
Value = (GetData (BubbaBaes + BUBBA_HIGH_DATA) << 8) |
         GetData (BubbaBaes + BUBBA_LOW_DATA)
```

Example 3-12. Horizontal Spacing Examples

3.11 Standard C Constructs

The following sections describe the code conventions for the following C constructs:

- Subroutines
- Function calls
- Boolean expressions
- Conditional expressions
- Loop expressions
- Switch expressions

3.11.1 Subroutines

Function prototype declarations and function implementations use the same coding conventions. The first line is the data type of the return value, left justified. If the function is a protocol member, then the next line is **EFIAPI**. The next line is the name of the function that follows the function naming conventions, followed by a space and an open parenthesis. This line is followed by the argument list with one argument per line, indented by two spaces. The type names for each argument are aligned on the same column, and the beginning of the field name for each argument is also aligned in the same column. The last argument is followed by a line with a close parenthesis that is also indented by two spaces.

Each argument type definition is preceded by an **IN** and/or **OUT** modifiers. The modifiers are used to indicate whether the argument is an input or output variable. The **IN** variables are listed first followed by the **OUT** variables. If data is both passed in and passed out through a variable, then it is marked as both **IN** and **OUT**. A buffer that is passed into a routine that modifies the contents of the buffer is marked as both **IN** and **OUT**. Table 3-5 shows the code conventions that are used for **IN** and **OUT**.

Table 3-5. IN and OUT Usage

Mnemonic	Description
IN	Passed by value. For C, this mnemonic is any argument whose name is not preceded by an asterisk (*). Placed at the beginning of a line containing a function argument.
IN	Passed by reference, and referenced data is not modified by the routine. An asterisk precedes the argument name. Placed at the beginning of a line containing a function argument.
OUT	Passed by reference, and the referenced data is modified by the routine. The passed-in state of the referenced data is not used by the routine. Placed at the beginning of a line containing a function argument.
IN OUT	Passed by reference, and the passed-in referenced data is consumed and then modified by the routine. Placed at the beginning of a line containing a function argument.
OPTIONAL	If argument is a **NULL** pointer, it is not present. If the value is not **NULL**, it is a valid argument. Placed at the end of a line containing a function argument.

Function implementations are then followed by a function comment header and the function body. The function body starts with the declaration of all the local variables that are used in the function. Local variable declarations may not be spread throughout the function. There is one local variable declared per line, indented by two spaces. Like the argument declarations, local variable names are aligned in the same column. Local variables declarations are not commented. Instead, they have self-describing names. If comments are required for complex local variable declarations, then the comments are placed in the include file that defines the complex data type, or the comments are placed in the function's comment header. Local variables are not initialized as part of their declaration. Instead, the local variables are initialized as part of the code body that follows the local variable declarations. Example 3-13 shows an example of the function prototype and function implementation.

```
//
// Function Prototype Declaration
//
EFI_STATUS
EFIAPI
FooName (
  IN     UINTN  Arg1,
  IN     UINTN  Arg2, OPTIONAL
  OUT    UINTN  *Arg3,
  IN OUT UINTN  *Arg4
  );

//
// Function Implementaion
//
EFI_STATUS
EFIAPI
FooName (
  IN     UINTN  Arg1,
  IN     UINTN  Arg2, OPTIONAL
  OUT    UINTN  *Arg3,
  IN OUT UINTN  *Arg4
  )
/*++

Routine Description:

  <<description>>

Arguments:

  <<argument names and purposes>>

Returns:

  <<description of possible return values>>

--*/
{
  UINTN  LocalOne;
  UINTN  LocalTwo;
  UINTN  LocalThree;
  . . .
}
```

Example 3-13. C Subroutine

3.11.2 Calling Functions

Function calls contain the function name, followed by a space and an open parenthesis. If all the arguments fit on one line, then the arguments are separated by commas on that line. If the arguments do not all fit on the same line, then each argument is placed on its own line, indented two spaces from the first character of the function being called. Example 3-14 shows the code conventions for calling functions.

```
//
// Function and arguments fit on a single line
//
Foo (A, B, C);

//
// Function and arguments do not fit on a single line
//
Status = BlockIo->ReadBlocks (
                    BlockIo,
                    MediaId,
                    Lba,
                    BufferSize,
                    Buffer
                    );
```

Example 3-14. Calling Functions

3.11.3 Boolean Expressions

Boolean tests are constructed using the following conventions. Boolean values and variables of type **BOOLEAN** do not require explicit comparisons to **TRUE** or **FALSE**. Non-Boolean values or variables use comparison operators (**==**, **!=**, **>**, **<**, **>=**, **<=**). A comparison of any pointer to zero is done with the **NULL** constant. Example 3-15 shows an example of the code conventions for Boolean expressions.

```
BOOLEAN  Done;
UINTN    Index;
VOID     *Ptr;

//
//          Incorrect                        Correct
//   =======================            =========================
     if (Index) {                       if (Index != 0) {
     if (!Index) {                      if (Index == 0) {
     if (Done == TRUE) {                if (Done) {
     if (Done == FALSE) {               if (!Done) {
     if (Ptr) {                         if (Ptr != NULL) {
     if (Ptr == 0) {                    if (Ptr == NULL) {
```

Example 3-15. Boolean Expressions

3.11.4 Conditional Expressions

Example 3-16 contains examples of the coding conventions for condition expressions, including the following:

- An **if** construct
- An **if**/**else** construct
- An **if**/**else if**/**else** construct
- A nested **if** construct

The **if** statement is followed by an open brace (**{**) even if the body contains only one statement. The body is indented two spaces, and the close brace (**}**) is placed on its own line and shares the same indentation as the **if** statement. The close brace (**}**) may optionally be followed by an **else if** or an **else** statement and an additional open brace (**{**). The body for the **else if** and the **else** statements are also indented by two spaces.

```
//
// IF construct
//
if (A > B) {
  IamTheCode ();
}

//
// IF / ELSE construct
//
if (Pointer == NULL) {
  IamTheCode ();
} else {
  IamTheCode ();
}

//
// IF / ELSE IF / ELSE construct
//
if (Done == TRUE) {
  IamTheCode();
} else if (A < B) {
  IamTheCode();
} else {
  IamTheCode();
}

//
// Nested IF construct
//
if (A > 10) {
  if (A > 20) {
    IamTheCode ();
  } else {
    if (A == 15) {
      IamTheCode ();
    } else {
      IamTheCode ();
    }
  }
}
```

Example 3-16. Conditional Expressions

3.11.5 Loop Expressions

Example 3-17 contains examples of the coding conventions for loop expressions, including the following:

- A **while** loop
- A **do** loop
- A **for** loop
- A nested **for** loop

The loop statements are followed by an open brace (**{**) on the same line even if the body contains only one statement. The body is indented two spaces, and the close brace (**}**) is placed on its own line and shares the same indentation as the loop statement.

```
//
// WHILE Loop construct
//
while (Pointer != NULL) {
  IamTheCode();
}

//
// DO Loop construct
//
do {
  IamTheCode();
} while (A < B);

//
// FOR Loop construct
//
for (Index = 0; Index < MAX_INDEX; Index++) {
  IamTheCode(Index);
}

//
// Nested Loop construct
//
for (Column = 0; Column < MAX_COLUMN; Column++) {
  for (Ror = 0; Row < MAX_ROW; Row++) {
    IamTheCode(Index);
  }
}
```

Example 3-17. Loop Expressions

3.11.6 Switch Expressions

Example 3-18 contains an example of a **switch** statement. The **switch** statement is followed by an open brace (**{**) on the same line. The **case** statements and **default** statements are placed on their own lines at the same indentation level as the **switch** statement. The body of the **case** statements and **default** statements are indented two spaces, and the close brace (**}**) is placed on its own line and shares the same indentation as the **switch** statement.

```
switch (Variable) {

case 1:
  IamTheCode ();
  break;

case 2:
  IamTheCode ();
  break;

default:
  IamTheCode ();
  break;
}
```

Example 3-18. Switch Expressions

3.11.7 Goto Expressions

In general, the **goto** construct is not used. However, it is acceptable to use **goto** for error handling and thus exiting a routine in an error case. A **goto** allows the error exit code to be contained in one place at the end of a routine. This location reduces software life cycle maintenance issues, as there can be one copy of error cleanup code per routine. If a **goto** is not used for this case, then the error cleanup code tends to get replicated multiple times, which tends to induce errors in code maintenance. The **goto** follows the normal rules for C code. The labels are left justified. Example 3-19 shows the code convention for **goto** expressions.

```
{
  EFI_STATUS  Status;

  . . .
  Status = IAmTheCode ();
  if (EFI_ERROR (Status)) {
    goto ErrorExit;
  }

  IDoTheWork ();

ErrorExit:
  . . .
  return Status;
}
```

Example 3-19. Goto Expressions

3.12 EFI File Templates

This section contains templates and guidelines for creating files for EFI protocols, EFI GUIDs, and EFI drivers. The naming conventions for the driver entry point and the functions exported by a driver that are presented here will guarantee that a unique name is produced for every function, which aides in call stack analysis when root-causing driver issues. The following expressions are used throughout this section to show where protocol names, GUID names, function names, and driver names should be substituted in a file template:

<<ProtocolName>>	Represents the name of a protocol that follows the function or variable naming convention, which capitalizes only the first letter of each word (e.g., **DiskIo**).
<<PROTOCOL_NAME>>	Represents the name of a protocol that follows the data structure naming convention, which capitalizes all the letters and separates each word with an underscore '_' (e.g., **DISK_IO**).
<<FunctionNameN>>	Represents the *n*th name of the protocol member functions that follow the function or variable naming convention, which capitalizes only the first letter of each word (e.g., **ReadDisk**).
<<FUNCTION_NAMEn>>	Represents the *n*th name of the protocol member functions that follows the data structure naming convention, which capitalizes all the letters and separates each word with an underscore '_' (e.g., **READ_DISK**).
<<GuidName>>	Represents the name of a GUID that follows the function or variable naming convention, which capitalizes only the first letter of each word (e.g., **GlobalVariable**).
<<GUID_NAME>>	Represents the name of a GUID that follows the data structure naming convention, which capitalizes all the letters and separates each word with an underscore '_' (e.g., **GLOBAL_VARIABLE**).
<<DriverName>>	Represents the name of a driver that follows the function or variable naming convention, which capitalizes only the first letter of each word (e.g., **Ps2Keyboard**).
<<DRIVER_NAME>>	Represents the name of a driver that follows the data structure naming convention, which capitalizes all the letters and separates each word with an underscore '_' (e.g., **PS2_KEYBOARD**).
<<DriverVersion>>	Value that represents the version of the driver. Values from 0x0–0x0f and 0xfffffff0–0xffffffff are reserved for EFI drivers that are written by OEMs for integrated devices. Values from 0x10–0xffffffef are reserved for EFI drivers that are written by IHVs.

<<ProtocolNameCn>>	Represents the *n*th name of a protocol that is consumed by an EFI driver and follows the function or variable naming convention, which capitalizes only the first letter of each word (e.g., **DiskIo**).
<<PROTOCOL_NAME_CN>>	Represents the *n*th name of a protocol that is consumed by an EFI driver and follows the data structure naming convention, which capitalizes all the letters and separates each word with an underscore '_' (e.g., **DISK_IO**).
<<ProtocolNamePm>>	Represents the *m*th name of a protocol that is produced by an EFI driver and follows the function or variable naming convention, which capitalizes only the first letter of each word (e.g., **DiskIo**).
<<PROTOCOL_NAME_PM>>	Represents the *m*th name of a protocol that is produced by an EFI driver and follows the data structure naming convention, which capitalizes all the letters and separates each word with an underscore '_' (e.g., **DISK_IO**).

3.12.1 Protocol File Templates

Protocols are placed in the directory **\Efi1.1\Edk\Protocol\<<ProtocolName>>**. These directories contain a **<<ProtocolName>>.c** and a **<<ProtocolName>>.h** file. The **<<ProtocolName>>.h** file has the contents shown in Example 3-20. The file header comments and function header comments have been omitted from this template. The GUID shown is an illegal GUID that is composed of all zeros. Every new protocol must generate a new GUID. **GUIDGEN**, which is shipped with Microsoft Visual Studio, can be used to generate new GUIDs. The **<<ProtocolName>>.c** file has the contents shown in Example 3-21. The file header comments have also been removed from this template. The full source to the Disk I/O Protocol is included in Appendix D for reference.

```
#ifndef __EFI_<<PROTOCOL_NAME>>_H__
#define __EFI_<<PROTOCOL_NAME>>_H__

#define EFI_<<PROTOCOL_NAME>>_PROTOCOL_GUID \
  { 0x00000000, 0x0000, 0x0000, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00 }

EFI_INTERFACE_DECL (_EFI_<<PROTOCOL_NAME>>_PROTOCOL);

typedef
EFI_STATUS
(EFIAPI *EFI_<PROTOCOL_NAME>>_<<PROTOCOL_FUNCTION_NAME1>>) (
  IN EFI_<<PROTOCOL_NAME>>_PROTOCOL *This,
  //
  // Place additional function arguments here.
  //
  );

typedef
EFI_STATUS
(EFIAPI *EFI_<PROTOCOL_NAME>>_<<PROTOCOL_FUNCTION_NAME2>>) (
  IN EFI_<<PROTOCOL_NAME>>_PROTOCOL *This,
  //
```

```
  // Place additional function arguments here.
  //
  );

// . . .

typedef
EFI_STATUS
(EFIAPI *EFI_<PROTOCOL_NAME>>_<<PROTOCOL_FUNCTION_NAMEn>>) (
  IN EFI_<<PROTOCOL_NAME>>_PROTOCOL *This,
  //
  // Place additional function arguments here.
  //
  );

typedef struct _EFI_<<PROTOCOL_NAME>>_PROTOCOL {
  EFI_<<PROTOCOL_NAME>>_<<FUNCTION_NAME1>>  <<FunctionName1>>;
  EFI_<<PROTOCOL_NAME>>_<<FUNCTION_NAME2>>  <<FunctionName2>>;
  // . . .
  EFI_<<PROTOCOL_NAME>>_<<FUNCTION_NAMEn>>  <<FunctionNameN>>;
  //
  // Place protocol data fields here
  //
} EFI_<<PROTOCOL_NAME>>_PROTOCOL;

extern EFI_GUID gEfi<<ProtocolName>>ProtocolGuid;

#endif
```

Example 3-20. Protocol Include File

```
#include "Efi.h"
#include EFI_PROTOCOL_DEFINITION (<<ProtocolName>>)

EFI_GUID gEfi<<ProtocolName>>ProtocolGuid = EFI_<<PROTOCOL_NAME>>_PROTOCOL_GUID;

EFI_GUID_STRING (&gEfi<<ProtocolName>>ProtocolGuid, "ShortString",
"LongString");
```

Example 3-21. Protocol C File

3.12.2 GUID File Templates

GUIDs and their associated data structures are declared just like protocols. The only difference is that GUIDs are placed in the directory **\Efi1.1\Edk\Guid\<<GuidName>>**. These directories contain a **<<GuidName>>.c** and a **<<GuidName>>.h** file. The **<<GuidName>>.h** file has the contents shown in Example 3-22. The file header comments have been omitted from this template. The GUID shown is an illegal GUID that is composed of all zeros. The **GUIDGEN** tool can be used to generate new GUIDs. The **<<GuidName>>.c** file has the contents shown in Example 3-23. The file header comments have also been omitted from this template. The full source to the EFI global variable GUID is included in Appendix D for reference. GUIDs are added to the **\Efi1.1\Edk\Guid\<<GuidName>>** directory when a GUID is required by more than one EFI component. If a GUID is required only by a single EFI driver, then it can be declared with the source code to the EFI driver.

```
#ifndef __<<GUID_NAME>>_H__
#define __<<GUID_NAME>>_H__

#define EFI_<<GUID_NAME>>_GUID \
  { 0x00000000, 0x0000, 0x0000, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00 }

typedef struct {
  //
  // Place GUID specific data fields here
  //
} EFI_<<GUID_NAME>>_GUID;

extern EFI_GUID gEfi<<GuidName>>Guid;

#endif
```

Example 3-22. GUID Include File

```
#include "Efi.h"
#include EFI_GUID_DEFINITION (<<GuidName>>)

EFI_GUID gEfi<<GuidName>>Guid = EFI_<<GUID_NAME>>_GUID;

EFI_GUID_STRING (&gEfi<<GuidName>>Guid, "", "");
```

Example 3-23. GUID C File

3.12.3 Including a Protocol or a GUID

Any code that produces or consumes a protocol must include the protocol definitions, and any code that produces or consumes a GUID must include the GUID definitions. A protocol can be included by using the **EFI_PROTOCOL_DEFINITION ()** macro. However, it is recommended that aliases of this macro be used that advertise if the protocol is being consumed or produced. The **EFI_PROTOCOL_CONSUMER ()** macro is used by EFI drivers to include protocol definitions that are consumed, and the **EFI_PROTOCOL_PRODUCER ()** macro is used by EFI drivers to include protocol definitions that are produced. There is no difference between these two macros in function, but they could potentially be used by source code analysis tools in the future to examine the relationships between various EFI drivers. A GUID is included by using the

EFI_GUID_DEFINITION () macro. Example 3-24 shows the different macros that can be used to include the Block I/O Protocol and the EFI global variable GUID.

```
#include EFI_PROTOCOL_DEFINITION (BlockIo)
#include EFI_PROTOCOL_PRODUCER (BlockIo)
#include EFI_PROTOCOL_CONSUMER (BlockIo)
#include EFI_GUID_DEFINITION (GlobalVariable)
```

Example 3-24. Including a Protocol or a GUID

3.12.4 EFI Driver Template

EFI drivers are placed below the **\Efi1.1\Edk\Drivers** directory. The directory structure in this area does not have to be flat. Closely related drivers may be placed in subdirectories. The directory name for an EFI driver is typically of the form **<<DriverName>>**. For example, the PS/2* keyboard driver is in the **\Efi1.1\Edk\Drivers\Ps2Keyboard** directory.

Simple EFI drivers will typically have the following two files in their driver directory:

- **<<DriverName>>.h**
- **<<DriverName>>.c**

The **<<DriverName>>.h** file includes the standard EFI include files, the EFI Driver Library declarations, and any protocol or GUID files that the driver either consumes or produces. In addition, the **<<DriverName>>.h** file contains the function prototypes of all the public APIs that are produced by the driver. The **<<DriverName>>.c** file contains the driver entry point. If an EFI driver produces the Driver Binding Protocol, then the **<<DriverName>>.c** file typically contains the **Supported()**, **Start()**, and **Stop()** services. The **<<DriverName>>.c** file may also contain the services for the protocol(s) that the EFI driver produces. Complex EFI drivers that produce more than one protocol may be broken up into multiple source files. The natural organization is to place the implementation of each protocol that is produced in a separate file of the form **<<ProtocolName>>.c** or **<<DriverName>><<ProtocolName>>.c**. For example, the disk I/O driver produces the Driver Binding Protocol, the Disk I/O Protocol, and the Component Name Protocol. The **DiskIo.c** file contains the Driver Binding Protocol and Disk I/O Protocol implementations. The **ComponentName.c** file contains the implementation of the Component Name Protocol.

3.12.5 <<DriverName>>.h File

Example 3-25 below shows the main include file for an EFI driver that consumes ***n*** protocols and produces ***m*** protocols. The file header comments and the function prototypes for the exported APIs have been omitted in this template. The full source to the include file for the disk I/O driver is included in Appendix D for reference. An EFI driver include file contains the following:

- **#ifndef** / **#define** for the driver include file
- **#include** statements for the standard EFI and EFI Driver Library include files.
- **#include** statements for all the protocols and GUIDs that are consumed by the driver.
- **#include** statements for all the protocols and GUIDs that are produced by the driver.
- **#define** for a unique signature that is used in the private context data structure (see chapter 8).

- **typedef struct** for the private context data structure (see chapter 8).
- **#define** statements to retrieve the private context data structure from each protocol that is produced (see chapter 8).
- **extern** statements for the global variables that the driver produces.
- Function prototype for the driver's entry point.
- Function prototypes for all of the APIs in the produced protocols
- **#endif** statement for the driver include file

```
#ifndef __EFI_<<DRIVER_NAME>>_H__
#define __EFI_<<DRIVER_NAME>>_H__

#include "Efi.h"
#include "EfiDriverLib.h"

//
// Driver Consumed Protocol Prototypes
//
#include EFI_PROTOCOL_CONSUMER (<<ProtocolNameC1>>)
#include EFI_PROTOCOL_CONSUMER (<<ProtocolNameC2>>)
// . . .
#include EFI_PROTOCOL_CONSUMER (<<ProtocolNameCn>>)

//
// Driver Produced Protocol Prototypes
//
#include EFI_PROTOCOL_PRODUCER (<<ProtocolNameP1>>)
#include EFI_PROTOCOL_PRODUCER (<<ProtocolNameP2>>)
// . . .
#include EFI_PROTOCOL_PRODUCER (<<ProtocolNamePm>>)

//
// Private Data Structure
//
#define <<DRIVER_NAME>>_PRIVATE_DATA_SIGNATURE  EFI_SIGNATURE_32
('A','B','C','D')

typedef struct {
  UINTN                   Signature;
  EFI_HANDLE              Handle;

  //
  // Pointers to consumed protocols
  //
  EFI_<<PROTOCOL_NAME_C1>>_PROTOCOL  *<<ProtocolNameC1>>;
  EFI_<<PROTOCOL_NAME_C2>>_PROTOCOL  *<<ProtocolNameC2>>;
  // . . .
  EFI_<<PROTOCOL_NAME_Cn>>_PROTOCOL  *<<ProtocolNameCn>>;

  //
  // Produced protocols
  //
  EFI_<<PROTOCOL_NAME_P1>>_PROTOCOL  <<ProtocolNameP1>>;
  EFI_<<PROTOCOL_NAME_P2>>_PROTOCOL  <<ProtocolNameP2>>;
  // . . .
  EFI_<<PROTOCOL_NAME_Pm>>_PROTOCOL  <<ProtocolNamePm>>;

  //
```

```
  // Private functions and data fields
  //
} <<DRIVER_NAME>>_PRIVATE_DATA;

#define <<DRIVER_NAME>_PRIVATE_DATA_FROM_<<PROTOCOL_NAME_P1>>_THIS(a) \
  CR(                                                                 \
    a,                                                                \
    <<DRIVER_NAME>>_PRIVATE_DATA,                                     \
    <<ProtocolNameP1>>,                                               \
    <<DRIVER_NAME>>_PRIVATE_DATA_SIGNATURE                            \
  )

#define <<DRIVER_NAME>_PRIVATE_DATA_FROM_<<PROTOCOL_NAME_P2>>_THIS(a) \
  CR(                                                                 \
    a,                                                                \
    <<DRIVER_NAME>>_PRIVATE_DATA,                                     \
    <<ProtocolNameP2>>,                                               \
    <<DRIVER_NAME>>_PRIVATE_DATA_SIGNATURE                            \
  )

// . . .

#define <<DRIVER_NAME>_PRIVATE_DATA_FROM_<<PROTOCOL_NAME_Pm>>_THIS(a) \
  CR(                                                                 \
    a,                                                                \
    <<DRIVER_NAME>>_PRIVATE_DATA,                                     \
    <<ProtocolNamePm>>,                                               \
    <<DRIVER_NAME>>_PRIVATE_DATA_SIGNATURE                            \
  )

//
// Required Global Variables
//
extern EFI_DRIVER_BINDING_PROTOCOL       g<<DriverName>>DriverBinding;

//
// Optional Global Variables
//
extern EFI_COMPONENT_NAME_PROTOCOL       g<DriverName>>ComponentName;
extern EFI_DRIVER_CONFIGURATION_PROTOCOL g<DriverName>>DriverConfiguration;
extern EFI_DRIVER_DIAGNOSTICS_PROTOCOL   g<DriverName>>DriverDiagnostics;

//
// Function prototype for the driver's entry point
//
EFI_STATUS
EFIAPI
<<DriverName>>DriverEntryPoint (
  IN EFI_HANDLE        ImageHandle,
  IN EFI_SYSTEM_TABLE  *SystemTable
  );

//
// Function ptototypes for the APIs in the Produced Protocols
//

#endif
```

Example 3-25. Driver Include File Template

3.12.6 <<DriverName>>.c File

The template in Example 3-26 below shows the main source file for an EFI driver. The file header comments and function header comments have been omitted in this template. This template contains empty protocol functions. The functions from the various protocols that an EFI driver may produce are discussed in later chapters. The full source to the include file for the disk I/O driver is included in Appendix D for reference. An EFI source file contains the following:

- **#include** statement for **<<DriverName>>.h**.
- Global variable declarations
- Declaration of the EFI driver's entry point function
- The EFI driver entry point function
- The **Supported()**, **Start()**, and **Stop()** functions
- Implementation of the APIs from the produced protocols

```
#include "<<DriverName>>.h"

EFI_DRIVER_BINDING_PROTOCOL g<<DriverName>>DriverBinding = {
  <<DriverName>>DriverBindingSupported,
  <<DriverName>>DriverBindingStart,
  <<DriverName>>DriverBindingStop,
  <<DriverVersion>>,
  NULL,
  NULL
};

EFI_DRIVER_ENTRY_POINT (<<DriverName>>DriverEntryPoint)

EFI_STATUS
EFIAPI
<<DriverName>>DriverEntryPoint (
  IN EFI_HANDLE        ImageHandle,
  IN EFI_SYSTEM_TABLE  *SystemTable
  )
{
  return EfiLibInstallAllDriverProtocols (
           ImageHandle,
           SystemTable,
           &g<<DriverName>>DriverBinding,
           ImageHandle,
           &g<<DriverName>>ComponentName,
           &g<<DriverName>>DriverConfiguration,
           &g<<DriverName>>DriverDiagnostics
           );
}

EFI_STATUS
EFIAPI
<<DriverName>>DriverBindingSupported (
  IN EFI_DRIVER_BINDING_PROTOCOL  *This,
  IN EFI_HANDLE                   ControllerHandle,
  IN EFI_DEVICE_PATH_PROTOCOL     *RemainingDevicePath  OPTIONAL
  )
{
}
```

```
EFI_STATUS
EFIAPI
<<DriverName>>DriverBindingStart (
  IN EFI_DRIVER_BINDING_PROTOCOL  *This,
  IN EFI_HANDLE                   ControllerHandle,
  IN EFI_DEVICE_PATH_PROTOCOL     *RemainingDevicePath  OPTIONAL
  )
{
}

EFI_STATUS
EFIAPI
<<DriverName>>DriverBindingStop (
  IN  EFI_DRIVER_BINDING_PROTOCOL  *This,
  IN  EFI_HANDLE                   ControllerHandle,
  IN  UINTN                        NumberOfChildren,
  IN  EFI_HANDLE                   *ChildHandleBuffer
  )
{
}

//
// Implementations of the APIs in the produced protocols
// The following template is for the mth function of the nth protocol produced
// It also shows how to retrieve the private context structure from this arg
//
EFI_STATUS
EFIAPI
<<DriverName>><<ProtocolNamePn>><<FunctionNameM>> (
  IN EFI_<<PROTOCOL_NAME_PN>>_PROTOCOL *This,
  //
  // Additional function arguments here.
  //
  )
{
  <<DRIVER_NAME>>_PRIVATE_DATA  *Private;

  //
  // Use This pointer to retrieve the private context structure
  //
  Private = <<DRIVER_NAME>_PRIVATE_DATA_FROM_<<PROTOCOL_NAME_Pn>>_THIS (This);
}
```

Example 3-26. Driver Implementation Template

3.12.7 <<ProtocolName>>.c or <<DriverName>><<ProtocolName>>.c File

More complex EFI drivers may break the implementation into several source files. The natural boundary is to implement one protocol per file. The template in Example 3-27 below shows the main source file for one protocol of an EFI driver. The file header comments and function header comments have been omitted in this template. This template shows only empty protocol functions. The functions from the various protocols that an EFI driver may produce are discussed in later chapters. The full source to the Component Name Protocol for the disk I/O driver is included in Appendix D for reference. An EFI protocol source file contains the following:

- **`#include`** statement for **`<<DriverName>>.h`**..
- Global variable declaration. This declaration applies only to protocols such as the Component Name, Driver Configuration, and Driver Diagnostics Protocols. Protocols that produce I/O services should never be declared as a global variable. Instead, they are declared in the private context structure that is dynamically allocated in the **`Start()`** function (see chapter 8).
- Implementation of the APIs from the produced protocols.

```
#include "<<DriverName>>.h"

//
// Protocol Global Variables
//
EFI_<<PROTOCOL_NAME_PN>>_PROTOCOL g<<DriverName>><<ProtocolNamePn>> = {
  // . . .
};

//
// Implementations of the APIs in the produced protocols
// The following template is for the mth function of the nth protocol produced
// It also shows how to retrieve the private context structure from the This arg
//
EFI_STATUS
EFIAPI
<<DriverName>><<ProtocolNamePn>><<FunctionName1M>> (
  IN EFI_<<PROTOCOL_NAME_PN>>_PROTOCOL *This,
  //
  // Additional function arguments here.
  //
  )
{
  <<DRIVER_NAME>>_PRIVATE_DATA  Private;

  //
  // Use This pointer to retrieve the private context structure
  //
  Private = <<DRIVER_NAME>_PRIVATE_DATA_FROM_<<PROTOCOL_NAME_Pn>>_THIS (This);
}
```

Example 3-27. Protocol Implementation Template

3.12.8 EFI Driver Library

Most EFI drivers use the EFI Driver Library because it simplifies the implementation and reduces the size of both the source code and the executable. All EFI drivers that use the EFI Driver Library must include the **EfiDriverLib.h** file. There are three different library initialization functions that are available from the EFI Driver Library. All three of these functions initialize the global variables **gST**, **gBS**, and **gRT**. These three global variables are pointers to the EFI System Table, the EFI Boot Services Table, and the EFI Runtime Services Table, respectively. There are also a number of function and macros available from the EFI Driver Library. They are all documented in the *EFI Driver Library Specification*. The code fragment in Example 3-28 below shows examples of the different EFI Driver Library initialization functions. The first example is typically used by services drivers that do not follow the EFI Driver Model. The second example is typically used by simple EFI drivers that follow the EFI Driver Model and do not produce the optional driver-related protocols. The last example is typically used by more complex EFI drivers that follow the EFI Driver Model and produce at least one of the optional driver-related protocols.

```
//
// Initialize a service driver
//
EfiInitializeDriverLib (ImageHandle, SystemTable);

//
// Initialize a simple EFI driver that follows the EFI Driver Model
//
EfiLibInstallDriverBinding (
  ImageHandle,
  SystemTable,
  gDiskIoDriverBinding,
  ImageHandle
  );

//
// Initialize a complex EFI driver that follows the EFI Driver Model
//
EfiLibInstallAllDriverProtocols (
  ImageHandle,
  SystemTable,
  gDiskIoDriverBinding,
  ImageHandle,
  gDiskIoComponentName,
  gDiskIoDriverConfiguration,
  gDiskIoDriverDiagnostics
  );
```

Example 3-28. Initializing the EFI Driver Library

4
EFI Services

An EFI driver has a number of EFI Boot Services and EFI Runtime Services that are available to perform its functions. These services can be grouped into the following three areas:

- Commonly used services
- Rarely used services
- Services that should not be used from an EFI driver

Table C-1, Table C-2, and Table C-3 in Appendix C list each of these categories. The full function prototypes and descriptions for each of the services and their arguments described in this chapter are available in chapters 5 and 6 of the *EFI 1.10 Specification*.

4.1 Services That EFI Drivers Commonly Use

Table 4-1 contains the list of EFI services that are commonly used by EFI drivers. The following sections provide a brief description of each service along with code examples on how they are typically used by EFI drivers. This table is also present in Appendix C for reference purposes.

Table 4-1. EFI Services That Are Commonly Used by EFI Drivers

Type	Service	Type	Service
BS	gBS->AllocatePool()	BS	gBS->InstallMultipleProtocolInterfaces()
BS	gBS->FreePool()	BS	gBS->UninstallMultipleProtocolInterfaces()
BS	gBS->AllocatePages()	BS	gBS->LocateHandleBuffer()
BS	gBS->FreePages()	BS	gBS->LocateProtocol()
BS	gBS->SetMem()	BS	gBS->OpenProtocol()
BS	gBS->CopyMem()	BS	gBS->CloseProtocol()
		BS	gBS->OpenProtocolInformation()
BS	gBS->CreateEvent()		
BS	gBS->CloseEvent()	BS	gBS->RaiseTPL()
BS	gBS->SignalEvent()	BS	gBS->RestoreTPL()
BS	gBS->SetTimer()		
BS	gBS->CheckEvent()	BS	gBS->Stall()

4.1.1 Memory Services

4.1.1.1 gBS->AllocatePool() and gBS->FreePool()

These services are used by EFI drivers to allocate and free small buffers that are guaranteed to be aligned on an 8-byte boundary. These services are ideal for allocating and freeing data structures. Because the allocated buffers are guaranteed to be on an 8-byte boundary, an alignment fault will not be generated on Itanium-based platforms as long as all the fields of the data structure are natural aligned (not packed).

When a buffer is allocated on an IA-32 platform, the buffer will be allocated below 4 GB, so it is guaranteed to be accessible by an IA-32 processor that is executing in flat physical mode. When a buffer is allocated on an Itanium-based platform, the buffer will be allocated somewhere in the 64-bit address space that is available to the Itanium processor in physical mode. This location means that buffers above 4 GB may be allocated on Itanium-based platforms if there is system memory above 4 GB. It is important to note that care must be taken when pointers are converted on Itanium-based platforms. All EFI drivers must be aware that pointers may contain values above 4 GB, and care must be taken to never strip the upper address bits. If the upper address bits are stripped, then the driver will work on IA-32 systems and Itanium-based platforms with small memory configurations, but not on Itanium-based platforms with larger memory configurations.

EFI Boot Service drivers will typically allocate and free buffers of type **EfiBootServicesData**, and EFI runtime drivers will typically allocate and free buffers of type **EfiRuntimeServicesData**. Most drivers that follow the EFI Driver Model will allocate private context structures in their **Start()** function and will free them in their **Stop()** function. EFI drivers may also dynamically allocate and free buffers as different I/O operations are performed. To prevent memory leaks, every allocation operation must have a corresponding free operation. The code fragment in Example 4-1 shows how these services can be used to allocate and free a buffer for a data structure from **EfiBootServicesData** memory. In addition, it shows how the same allocations can be performed using services from the EFI Driver Library. The EFI Driver Library provides services that reduce the size of the driver and make the driver code easier to read and maintain.

```
EFI_STATUS      Status;
IDE_BLK_IO_DEV  *IdeBlkIoDevice;

//
// Allocate a buffer for a data structure
//
Status = gBS->AllocatePool (
                EfiBootServicesData,
                sizeof (IDE_BLK_IO_DEV),
                (VOID**)&IdeBlkIoDevice
                );
if (EFI_ERROR (Status)) {
  return Status;
}

//
// Allocate the same buffer using an EFI Library Function
//
IdeBlkIoDevice = EfiLibAllocatePool (sizeof (IDE_BLK_IO_DEV));
if (IdeBlkIoDevice == NULL) {
```

```
  return EFI_OUT_OF_RESOURCES;
}

//
// Allocate the same buffer using an EFI Library Function that also
// initializes the contents of the buffer with zeros
//
IdeBlkIoDevice = EfiLibAllocateZeroPool (sizeof (IDE_BLK_IO_DEV));
if (IdeBlkIoDevice == NULL) {
  return EFI_OUT_OF_RESOURCES;
}

//
// Free the allocated buffer
//
Status = gBS->FreePool (IdeBlkIoDevice);
if (EFI_ERROR (Status)) {
  return Status;
}
```

Example 4-1. Allocate and Free Pool

4.1.1.2 gBS->AllocatePages() and gBS->FreePages()

These services are used by EFI drivers to allocate and free larger buffers that are guaranteed to be aligned on a 4 KB boundary. These services allow buffers to be allocated at any available address, at specific addresses, or below a specific address. However, EFI drivers should not make any assumptions about the organization of system memory, so allocating from specific addresses or below specific addresses is strongly discouraged. As a result, buffers are typically allocated with an allocation type of **AllocateAnyPages**, and allocation types of **AllocateMaxAddress** and **AllocateAddress** are not used.

When an allocation type of **AllocateAnyPages** is used on IA-32 platforms, the buffer will be allocated below 4 GB, so it is guaranteed to be accessible by an IA-32 processor that is executing in flat physical mode. When an allocation type of **AllocateAnyPages** is used on Itanium-based platforms, the buffer will be allocated somewhere in the 64-bit address space that is available to the Itanium processor in physical mode. This location means that buffers above 4 GB may be allocated on Itanium-based platforms if there is system memory above 4 GB. It is important to note this possible allocation because care must be taken when the physical address is converted to a pointer on Itanium-based platforms. All EFI drivers must be aware that pointers may contain values above 4 GB, and care must be taken to never strip the upper address bits. If the upper address bits are stripped, then the driver will work on IA-32 systems and Itanium-based platforms with small memory configurations, but not on Itanium-based platforms with larger memory configurations.

EFI Boot Service drivers will typically allocate and free buffers of type **EfiBootServicesData**, and EFI runtime drivers will typically allocate and free buffers of type **EfiRuntimeServicesData**. To prevent memory leaks, every allocation operation must have a corresponding free operation. The code fragment in Example 4-2 shows how these services can be used to allocate and free a buffer for a data structure from **EfiBootServicesData** memory.

```
EFI_STATUS            Status;
EFI_PHYSICAL_ADDRESS  PhysicalBuffer;
UINTN                 Size;
VOID                  *Buffer;

//
// Allocate the number of pages to hold Size bytes and return in PhysicalBuffer
//
Status = gBS->AllocatePages(
                AllocateAnyPages,
                EfiBootServicesData,
                EFI_SIZE_TO_PAGES(Size),
                &PhysicalBuffer
                );
if (EFI_ERROR (Status)) {
  return Status;
}

//
// Convert the physical address to a pointer. This mechanism is the best one
// that works on all IA-32 and Itanium-based platforms.
//
Buffer = (VOID *)(UINTN)PhysicalBuffer;

//
// Free the allocated buffer
//
Status = gBS->FreePages (RomBuffer, EFI_SIZE_TO_PAGES(RomBarSize));
if (EFI_ERROR (Status)) {
  return Status;
}
```

Example 4-2. Allocate and Free Pages

4.1.1.3 gBS->SetMem()

This service is used to initialize the contents of a buffer with a specified value. EFI drivers most commonly use this service to zero an allocated buffer, but it can be used to fill a buffer with other values too. The code fragment in Example 4-3 shows the same example from Example 4-1, but it uses this service to zero the contents of the allocated buffer.

```
EFI_STATUS      Status;
IDE_BLK_IO_DEV  *IdeBlkIoDevice;

//
// Allocate a buffer for a data structure
//
Status = gBS->AllocatePool (
                EfiBootServicesData,
                sizeof (IDE_BLK_IO_DEV),
                (VOID**)&IdeBlkIoDevice
                );
if (EFI_ERROR (Status)) {
  return Status;
}

//
// Zero the contents of the allocated buffer
//
gBS->SetMem (IdeBlkIoDevice, sizeof (IDE_BLK_IO_DEV), 0);
```

Example 4-3. Allocate and Free Buffer

4.1.1.4 gBS->CopyMem()

This service copies a buffer from one location to another. This service handles both aligned and unaligned buffers, and it even handles the rare case when buffers are overlapping. In the overlapping case, the requirement is that the destination buffer on exit from this service matches the contents of the source buffer on entry to this service. The code fragment in Example 4-4 shows how this service is typically used.

```
EFI_STATUS    Status;
MY_STRUCTURE  *SourceStructure;
MY_STRUCTURE  *DestinationStructure;

//
// Allocate a buffer for the destination buffer
//
DestinationStructure = EfiLibAllocatePool (sizeof (MY_STRUCTURE));
if (DestinationStructure == NULL) {
  return EFI_OUT_OF_RESOURCES;
}

//
// Copy the source buffer to the destination buffer
//
gBS->CopyMem (DestinationBuffer, SourceBuffer, sizeof (MY_STRUCTURE));
```

Example 4-4. Allocate and Copy Buffer

4.1.2 Handle Database and Protocol Services

4.1.2.1 gBS->InstallMultipleProtocolInterfaces() and gBS->UninstallMultipleProtocolInterfaces()

These services are used to do the following:

- Add handles to the handle database.
- Remove handles from the handle database.
- Add protocols to an existing handle in the handle database.
- Remove protocols from an existing handle in the handle database.

A handle in the handle database contains one or more protocols. A handle is not allowed to have zero protocols, and a handle is not allowed to have more than one instance of the same protocol on the same handle. A handle is added to the handle database when the first protocol is added. A handle is removed from the handle database when the last protocol is removed from the handle. These services support adding and removing more than one protocol at a time. If one protocol fails to be added to a handle, then none of the protocols are added to the handle. If one protocol fails to be removed from a handle, then none of the protocols are removed from the handle. The protocols are represented by a pointer to a protocol GUID and a pointer to the protocol interface. These services will parse pairs of arguments until a **NULL** pointer for the protocol GUID parameter is encountered. It is recommended that these services be used instead of **gBS->InstallProtocolInterface()** and **gBS->UninstallProtocolInterface()** because it results in small executables with source code that is easier to read. In addition, the **gBS->InstallMultipleProtocolInterfaces()** service will check to see if the same device path is being installed onto more than one handle in the handle database. This operation is not legal, so this additional error checking was added to section 5.3.1 of the *EFI 1.10 Specification*.

EFI drivers are required to add the **EFI_DRIVER_BINDING_PROTOCOL** to the handle database in their driver entry point. They may also add the following in the driver entry point:

- **EFI_DRIVER_CONFIGURATION_PROTOCOL**
- **EFI_DRIVER_DIAGNOSTICS_PROTOCOL**
- **EFI_COMPONENT_NAME_PROTOCOL**

These protocols are optional for some platforms and required for others. For example, *DIG64* requires these protocols for Itanium-based platforms. If an EFI driver is unloadable, then the protocols that were added in the driver entry point must be removed in the driver's **Unload()** function. The code fragment in Example 4-5 shows how these services would be used in a driver entry point and **Unload()** function.

```
EFI_DRIVER_BINDING_PROTOCOL  gMyDriverBinding = {
  MySupported,
  MyStart,
  MyStop,
  0x10,
  NULL,
  NULL
};

EFI_COMPONENT_NAME_PROTOCOL  gMyComponentName = {
```

```
  MyGetDriverName,
  MyGetControllerName,
  "eng"
};

EFI_HANDLE  ImageHandle;

//
// Install the Driver Binding Protocol and the Component Name Protocol
// onto the image handle that is passed into the driver entry point
//
Status = gBS->InstallMultipleProtocolInterfaces (
                &ImageHandle,
                &gEfiDriverBindingProtocolGuid, &gMyDriverBinding,
                &gEfiComponentNameProtocolGuid, &gMyComponentName,
                NULL
                );
if (EFI_ERROR (Status)) {
  return Status;
}

//
// Uninstall the Driver Binding Protocol and the Component Name Protocol
// from the handle that is passed into the Unload() function.
//
Status = gBS->UninstallMultipleProtocolInterfaces (
                ImageHandle,
                &gEfiDriverBindingProtocolGuid, &gMyDriverBinding,
                &gEfiComponentNameProtocolGuid, &gMyComponentName,
                NULL
                );
if (EFI_ERROR (Status)) {
  return Status;
}
```

Example 4-5. Install Driver Protocols

EFI device drivers will add protocols for I/O services to existing handles in the handle database in their **Start()** function and will remove those same protocols from those same handles in their **Stop()** function. EFI bus drivers may add protocols to existing handles, but they are also responsible for creating handles for the child device on that bus. This responsibility means that the EFI bus driver will typically add the **EFI_DEVICE_PATH_PROTOCOL** and an I/O abstraction for the bus type that the bus driver manages. For example, the PCI bus driver will create child handles with both the **EFI_DEVICE_PATH_PROTOCOL** and the **EFI_PCI_IO_PROTOCOL**. The bus driver may also optionally add the **EFI_BUS_SPECIFIC_DRIVER_OVERRIDE_PROTOCOL** to the child handles. The code fragment in Example 4-6 shows an example of a how a child handle can be created and additional protocols added and then destroyed.

```
EFI_HANDLE  ChildHandle;

//
// Add Device Path Protocol and Block I/O Protocol to a new handle
//
ChildHandle = NULL;
Status = gBS->InstallMultipleProtocolInterfaces (
                &ChildHandle,
                &gEfiDevicePathProtocol, DevicePath,
```

```
                &gEfiBlockIoProtocol,    BlockIo
                NULL
                );
if (EFI_ERROR (Status)) {
  return Status;
}

//
// Add the Disk I/O Protocol to the new handle created in the previous call
//
Status = gBS->InstallMultipleProtocolInterfaces (
                &ChildHandle,
                &gEfiDiskIoProtocol, DiskIo
                NULL
                );
if (EFI_ERROR (Status)) {
  return Status;
}

//
// Remove Device Path Protocol, Block I/O Protocol, and Disk I/O Protocol
// from the handle created above. Because this call will remove all the
// protocols from the handle, the handle will be removed from the handle
// database
//
Status = gBS->UnistallMultipleProtocolInterfaces (
                ChildHandle,
                &gEfiDevicePathProtocolGuid, DevicePath,
                &gEfiBlockIoProtocolGuid,    BlockIo,
                &gEfiDiskIoProtocolGuid,     DiskIo
                NULL
                );
if (EFI_ERROR (Status)) {
  return Status;
}
```

Example 4-6. Install I/O Protocols

Some EFI device drivers will add protocols to handles in the handle database with a **NULL** protocol interface. This case is known as a tag GUID because there are no data fields or services associated with the GUID. The code fragment in Example 4-7 shows an example of a how a tag GUID for hot-plug devices can be added and removed from a controller handle in the handle database.

```
EFI_HANDLE  ControllerHandle;

//
// Add the hot-plug device GUID to ControllerHandle
//
Status = gBS->InstallMultipleProtocolInterfaces (
                &ControllerHandle,
                &gEfiHotPlugDeviceGuid, NULL,
                NULL
                );
if (EFI_ERROR (Status)) {
  return Status;
}

//
// Remove the hot-plug device GUID from ControllerHandle
//
Status = gBS->UninstallMultipleProtocolInterfaces (
                ControllerHandle,
                &gEfiHotPlugDeviceGuid, NULL,
                NULL
                );
if (EFI_ERROR (Status)) {
  return Status;
}
```

Example 4-7. Install Tag GUID

NOTE

When an attempt is made to remove a protocol interface from a handle in the handle database, the EFI core firmware will check to see if any other EFI drivers are currently using the services of the protocol that is about to be removed. If EFI drivers are using that protocol interface, then the EFI core firmware will attempt to stop those EFI drivers with a call to **gBS->DisconnectController()**. *If the call to* **gBS->DisconnectController()** *fails, then the EFI core firmware will have to call* **gBS->ConnectController()** *to put the handle database back into the same state that it was in prior to the original call to* **gBS->UninstallMultipleProtocolInterfaces()**. *This call to* **gBS->ConnectController()** *has the potential to cause re-entrancy issues in EFI drivers that must be handled in the EFI driver. Please see chapter 21 for recommendations on how to test EFI drivers.*

4.1.2.2 gBS->LocateHandleBuffer()

This service retrieves a list of handles that meet a search criterion from the handle database. The following are the search options:

- Retrieve all the handles in the handle database.
- Retrieve the handles that support a specific protocol.
- Retrieve the handles that are in the notified state based on a prior **gBS->RegisterProtocolNotify()**.

It is not recommended that EFI drivers use **gBS->RegisterProtocolNotify()**, so only the first two cases will be covered here. The buffer that is returned by this service is allocated by the service, so an EFI driver that uses this service is responsible for freeing the returned buffer when the EFI driver no longer requires its contents. This service, along with **gBS->ProtocolsPerHandle()** and **gBS->OpenProtocolInformation()**, can be used to traverse the contents of the entire handle database. The algorithm for performing this traversal is in section 5.3.1 of the *EFI 1.10 Specification*.

The code fragment in Example 4-8 shows how all the handles in the handle database can be retrieved.

```
EFI_STATUS  Status;
UINTN       HandleCount;
EFI_HANDLE  *HandleBuffer;

//
// Retrieve the list of all the handles in the handle database. The number
// of handles in the handle database is returned in HandleCount, and the
// array of handle values is returned in HandleBuffer.
//
Status = gBS->LocateHandleBuffer (
                AllHandles,
                NULL,
                NULL,
                &HandleCount,
                &HandleBuffer
                );
if (EFI_ERROR (Status)) {
  return Status;
}

//
// Free the array of handles that was allocated by gBS->LocateHandleBuffer()
//
gBS->FreePool (HandleBuffer);
```

Example 4-8. Locate All Handles

The code fragment in Example 4-9 shows how all the handles that support the Block I/O Protocol can be retrieved and how the individual Block I/O Protocol instances can be retrieved using **gBS->OpenProtrocol()**.

```
EFI_STATUS              Status;
UINTN                   HandleCount;
EFI_HANDLE              *HandleBuffer;
UINTN                   Index;
EFI_BLOCK_IO_PROTOCOL   *BlockIo;

//
// Retrieve the list of handles that support the Block I/O Protocol from
// the handle database. The number of handles that support the Block I/O
// Protocol is returned in HandleCount, and the array of handle values is
// returned in HandleBuffer.
//
Status = gBS->LocateHandleBuffer (
                ByProtocol,
                &gEfiBlockIoProtocolGuid,
                NULL,
                &HandleCount,
                &HandleBuffer
                );
if (EFI_ERROR (Status)) {
  return Status;
}

//
// Loop through all the handles that support the Block I/O Protocol, and
// retrieve the instance of the Block I/O Protocol.
//
for (Index = 0; Index < HandleCount; Index++) {
  Status = gBS->OpenProtocol (
                  HandleBuffer[Index],
                  &gEfiBlockIoProtocolGuid,
                  (VOID **)&BlockIo,
                  ImageHandle,
                  NULL,
                  EFI_OPEN_PROTOCOL_GET_PROTOCOL
                  );
  if (EFI_ERROR (Status)) {
    return Status;
  }
  //
  // BlockIo can be used here to make block I/O service requests.
  //

}
//
// Free the array of handles that was allocated by gBS->LocateHandleBuffer()
//
gBS->FreePool (HandleBuffer);
```

Example 4-9. Locate Block I/O Protocol Handles

4.1.2.3 gBS->LocateProtocol()

This service finds the first instance of a protocol interface in the handle database. This service is typically used by EFI drivers to retrieve service protocols on service handles that are guaranteed to have at most one instance of the protocol in the handle database. The list of service protocols that are defined in the *EFI 1.10 Specification* includes the following:

- **EFI_PLATFORM_DRIVER_OVERRIDE_PROTOCOL**
- **EFI_UNICODE_COLLATION_PROTOCOL**
- **EFI_BIS_PROTOCOL**
- **EFI_DEBUG_SUPPORT_PROTOCOL**
- **EFI_DECOMPRESS_PROTOCOL**
- **EFI_EBC_PROTOCOL**

If there is a chance that more than one instance of a protocol may exist in the handle database, then **gBS->LocateHandleBuffer()** should be used. This service also supports retrieving protocols that have been notified with **gBS->RegisterProtocolNotify()**. It is not recommended that EFI drivers use **gBS->RegisterProtocolNotify()**, so this case will not be covered here. The code fragment in Example 4-10 shows how this service can be used to retrieve the **EFI_DECOMPRESS_PROTOCOL**.

```
EFI_STATUS                Status;
EFI_DECOMPRESS_PROTOCOL   *Decompress;

Status = gBS->LocateProtocol(
                &gEfiDecompressProtocolGuid,
                NULL,
                (VOID **)&Decompress
                );
if (EFI_ERROR (Status)) {
  return Status;
}
```

Example 4-10. Locate Decompress Protocol

4.1.2.4 gBS->OpenProtocol() and gBS->CloseProtocol()

These two services are used by EFI drivers to acquire and release the protocol interfaces that the EFI drivers require to produce their services. These services are some of the more complex services in EFI, but using them correctly is required for EFI drivers to produce their I/O abstractions and work well with other EFI drivers. EFI applications and EFI OS loaders can also use these services, but the discussion here will concentrate on the different ways that EFI drivers use these services.

gBS->OpenProtocol() is typically used by the **Supported()** and **Start()** functions of an EFI driver to retrieve a protocol interface that is installed on a handle in the handle database. A brief description of the parameters and return codes is presented here. A more complete description can be found in section 5.3.1 of the *EFI 1.10 Specification*. The function prototype of this service is shown below for reference.

```
typedef
EFI_STATUS
(EFIAPI *EFI_OPEN_PROTOCOL) (
  IN  EFI_HANDLE  Handle,
  IN  EFI_GUID    *Protocol,
  OUT VOID        **Interface OPTIONAL,
  IN  EFI_HANDLE  AgentHandle,
  IN  EFI_HANDLE  ControllerHandle,
  IN  UINT32      Attributes
  );
```

The only **OUT** parameter is the pointer to the protocol interface. All the other parameters are **IN** parameters that tell the EFI core why the protocol interface is being retrieved and by whom. The EFI core uses these **IN** parameters to track how each protocol interface is being used. This tracking information can be retrieved by using the **gBS->OpenProtocolInformation()** service. *AgentHandle* and *ControllerHandle* describe who is opening the protocol interface, and the *Attributes* parameter tells why the protocol interface is being opened. For EFI drivers, the *AgentHandle* parameter is typically the *DriverBindingHandle* field from the **EFI_DRIVER_BINDING_PROTOCOL**. Also, *Handle* and *ControllerHandle* are typically the same handle. The only exception is when a bus driver is opening a protocol on behalf of a child controller. The attributes that are used by EFI drivers are listed below. It is very important that EFI drivers use the correct attributes when a protocol interface is opened.

TEST_PROTOCOL Tests to see if a protocol interface is present on a handle. Typically used in the **Supported()** service of an EFI driver.

GET_PROTOCOL Retrieves a protocol interface from a handle.

BY_DRIVER Retrieves a protocol interface from a handle and marks that interface so it cannot be opened by other EFI drivers or EFI applications unless the other EFI driver agrees to release the protocol interface. This attribute is the most commonly used.

BY_DRIVER | EXCLUSIVE
Retrieves a protocol interface from a handle and marks the interface so it cannot be opened by other EFI drivers or EFI applications. This protocol will not be released until the driver that opened this attribute chooses to close it. This attribute is used very rarely.

BY_CHILD_CONTROLLER
Only used by bus drivers. A bus driver is required to open the parent I/O abstraction on behalf of each child controller that the bus driver produces. This requirement allows the EFI core to keep track of the parent/child relationships no matter how complex the bus hierarchies become.

The status code returned is also very important and must be examined by EFI drivers that use this service. The typical return codes are listed below.

EFI_SUCCESS The protocol interface was retrieved.

EFI_UNSUPPORTED The protocol interface is not present on the handle.

EFI_ALREADY_STARTED The EFI driver has already retrieved the protocol interface.

EFI_ACCESS_DENIED A different EFI driver has already retrieved the protocol interface.

EFI_INVALID_PARAMETER One of the parameters is invalid.

The **gBS->CloseProtocol()** service removes an element from the list of agents that are consuming a protocol interface. EFI drivers are required to close each protocol that they open, which is typically done in the **Stop()** function.

The code fragment in Example 4-11 shows the most common use of these services. This example shows a stripped-down version of a **Support()** function for a PCI device driver. It opens the PCI I/O Protocol **BY_DRIVER**, and then closes it if the open operation worked. If this driver wanted to open the PCI I/O Protocol exclusively, then an attribute of **BY_DRIVER | EXCLUSIVE** should be used. There are only a very few instances where **BY_DRIVER | EXCLUSIVE** should be used. These are cases where an EFI driver actually wants to gain exclusive access to a protocol, even if it requires stopping other EFI drivers to do so. This operation can be dangerous if the system requires the services produced by the EFI drivers that are stopped. One example is the debug port driver that opens the Serial I/O Protocol with the **BY_DRIVER | EXCLUSIVE** attribute. This attribute allows a debugger to take control of a serial port even if it is being used as a console device. If this device is the only console device in the system, then the user will lose the only console device when the debug port driver is started.

```
EFI_STATUS
XyzDriverBindingSupported (
  IN EFI_DRIVER_BINDING_PROTOCOL  *This,
  IN EFI_HANDLE                   ControllerHandle,
  IN EFI_DEVICE_PATH_PROTOCOL     *RemainingDevicePath
  )

{
  EFI_STATUS           Status;
  EFI_PCI_IO_PROTOCOL  *PciIo;

  Status = gBS->OpenProtocol (
                 ControllerHandle,
                 &gEfiPciIoProtocolGuid,
                 (VOID **)&PciIo,
                 This->DriverBindingHandle,
                 ControllerHandle,
                 EFI_OPEN_PROTOCOL_BY_DRIVER
                );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  gBS->CloseProtocol (
         ControllerHandle,
         &gEfiPciIoProtocolGuid,
         This->DriverBindingHandle,
```

```
          ControllerHandle
          );

  return Status;
}
```

Example 4-11. Open Protocol BY_DRIVER

The code fragment in Example 4-12 shows the same example as above, but it tests only for the presence of the PCI I/O Protocol using the **TEST_PROTOCOL** attribute. When **TEST_PROTOCOL** is used, the protocol does not have to be closed because a protocol interface is not returned when this open mode is used.

```
EFI_STATUS
XyzDriverBindingSupported (
  IN EFI_DRIVER_BINDING_PROTOCOL  *This,
  IN EFI_HANDLE                   ControllerHandle,
  IN EFI_DEVICE_PATH_PROTOCOL     *RemainingDevicePath
  )

{
  EFI_STATUS           Status;

  Status = gBS->OpenProtocol (
                 ControllerHandle,
                 &gEfiPciIoProtocolGuid,
                 NULL,
                 This->DriverBindingHandle,
                 ControllerHandle,
                 EFI_OPEN_PROTOCOL_TEST_PROTOCOL
                );
  return Status;
}
```

Example 4-12. Open Protocol by TEST_PROTOCOL

The code fragment in Example 4-13 shows the same example as above, but it retrieves the PCI I/O Protocol using the **GET_PROTOCOL** attribute. When **GET_PROTOCOL** is used, the protocol does not have to be closed.

NOTE

It can be dangerous to use this open mode because a protocol may be removed at any time, and a driver that uses **GET_PROTOCOL** *may attempt to use a stale protocol interface. There are a few places where it has to be used. An EFI driver should be designed to use* **BY_DRIVER** *as its first choice. However, there are cases where a different EFI driver has already opened the protocol that is required* **BY_DRIVER**. *The next best choice is to use* **GET_PROTOCOL**. *This scenario may occur when protocols are layered on top of each other, so that each layer uses the services of the layer immediately below. Each layer immediately below is opened* **BY_DRIVER**. *If a layer ever needs to skip around a layer to a lower-level service, then it is safe to use* **GET_PROTOCOL** *because the driver will be informed through the layers if the lower-level protocol is removed.*

The best example of this case in the *EFI Sample Implementation* is the FAT driver. The FAT driver uses the services of the Disk I/O Protocol to access the contents of the disk. However, the Disk I/O

Protocol does not have a flush service. Only the Block I/O Protocol has a flush service. The disk I/O driver opens the Block I/O Protocol **BY_DRIVER**, so the FAT driver is not allowed to also open the Block I/O Protocol **BY_DRIVER**. Instead, the FAT driver must use **GET_PROTOCOL**. This method is safe because the FAT driver will be indirectly notified if the Block I/O Protocol is removed when the Disk I/O Protocol is removed in response to the Block I/O Protocol being removed.

```
EFI_STATUS
XyzDriverBindingSupported (
  IN EFI_DRIVER_BINDING_PROTOCOL  *This,
  IN EFI_HANDLE                   ControllerHandle,
  IN EFI_DEVICE_PATH_PROTOCOL     *RemainingDevicePath
  )

{
  EFI_STATUS           Status;
  EFI_PCI_IO_PROTOCOL  *PciIo;

  Status = gBS->OpenProtocol (
                  ControllerHandle,
                  &gEfiPciIoProtocolGuid,
                  (VOID **)&PciIo,
                  This->DriverBindingHandle,
                  ControllerHandle,
                  EFI_OPEN_PROTOCOL_GET_PROTOCOL
                 );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Use the services of the PCI I/O Protocol here
  //
}
```

Example 4-13. Open Protocol by GET_PROTOCOL

The code fragment in Example 4-14 shows an example of a PCI bus driver opening the PCI Root Bridge I/O Protocol on behalf of a child PCI controller. This example shows the **BY_CHILD_CONROLLER** attribute being used. This attribute is typically used in the **Start()** function after the child handle has been created using **gBS->InstallMultipleProtocolInterfaces()**.

```
EFI_DRIVER_BINDING_PROTOCOL      gPciBusDriverBinding;

EFI_STATUS                       Status;
EFI_HANDLE                       ChildHandle;
EFI_DEVICE_PATH_PROTOCOL         *DevicePath;
EFI_PCI_IO_PROTOCOL              *PciIo;
EFI_HANDLE                       ControllerHandle;
EFI_PCI_ROOT_BRIDGE_IO_PROTOCOL  *PciRootBridgeIo;

ChildHandle = NULL;
Status = gBS->InstallMultipleProtocolInterfaces (
                &ChildHandle,
                &gEfiDevicePathProtocolGuid,  DevicePath,
                &gEfiPciIoProtocolGuid,       PciIo,
                NULL
```

```
                 );
if (EFI_ERROR (Status)) {
  return Status;
}

Status = gBS->OpenProtocol (
                ControllerHandle,
                &gEfiPciRootBridgeIoProtocolGuid,
                (VOID **)&PciRootBridgeIo,
                gPciBusDriverBinding.DriverBindingHandle,
                ChildHandle,
                EFI_OPEN_PROTOCOL_BY_CHILD_CONTROLLER
                );
if (EFI_ERROR (Status)) {
  return Status;
}
```

Example 4-14. Open Protocol BY_CHILD_CONTROLLER

4.1.2.5 gBS->OpenProtocolInformation()

The handle database contains the following:

- List of handles
- List of protocols on each handle
- List of agents that are currently using each protocol

This service retrieves the list of agents that are currently using a specific protocol interface that is registered in the handle database. An agent is an EFI driver or an EFI application that is using the services of a protocol interface. The **gBS->OpenProtocol()** service adds agents to the list, and the **gBS->CloseProtocol()** service removes agents from the list. An EFI driver will typically use this service to find the list of child handles that the EFI driver may have produced in previous calls to the **Start()**. The return buffer is allocated by the service, so the caller must free the buffer when the caller does not need the return buffer anymore. The code fragment in Example 4-15 retrieves the list of agents that are using the Serial I/O Protocol on a controller handle. This step is followed by a loop that counts the number of children that have been produced from the Serial I/O Protocol. The final step is to free the return buffer. This algorithm can be used by bus drivers that produce at most one child handle, and they need to check to see if a child handle has already been produced.

```
EFI_STATUS                          Status;
EFI_HANDLE                          ControllerHandle;
EFI_OPEN_PROTOCOL_INFORMATION_ENTRY *OpenInfo;
UINTN                               EntryCount;
UINTN                               Index;
UINTN                               NumberOfChildren;

Status = gBS->OpenProtocolInformation (
                ControllerHandle,
                &gEfiSerialIoProtocolGuid,
                &OpenInfo,
                &EntryCount
                );
if (EFI_ERROR (Status)) {
  return Status;
}
```

```
for (Index = 0, NumberOfChildren = 0; Index < EntryCount; Index++) {
  if (OpenInfo[Index].Attributes & EFI_OPEN_PROTOCOL_BY_CHILD_CONTROLLER) {
    NumberOfChildren++;
  }
}

gBS->FreePool (OpenInfoBuffer);
```

Example 4-15. Open Protocol Information

4.1.3 Task Priority Level Services

4.1.3.1 gBS->RaiseTPL() and gBS->RestoreTPL()

These two services are used to raise and restore the Task Priority Level (TPL) of the system. The most common use of these services is to implement a simple lock, or critical section, on global data structures. EFI runs on one thread on one processor. However, EFI does support a single timer interrupt. Because EFI code can run in interrupt context, it is possible that a global data structure can be accessed from both normal context and interrupt context. As a result, global data structures must be protected by a lock. The code fragment in Example 4-16 shows how these services can be used to implement a lock when the contents of a global variable are modified. The timer interrupt is blocked at **EFI_TPL_HIGH_LEVEL**, so most locks raise to this level.

```
UINT32  gCounter;

EFI_TPL  OldTpl;

//
// Raise the Task Priority Level to EFI_TPL_HIGH_LEVEL to block timer
// interrupts
//
OldTpl = gBS -> RaiseTPL(EFI_TPL_HIGH_LEVEL);

//
// Increment the global variable now that it is safe to do so.
//
gCounter++

//
// Restore the Task Priority Level to its original level
//
gBS -> RestoreTPL(OldTpl);
```

Example 4-16. Global Lock

These services are also used to raise the TPL during a blocking I/O transaction. Most drivers are required to raise the TPL to **EFI_TPL_NOTIFY** during blocking I/O operations. See Example 4-19 below for an example of a keyboard driver that raises the TPL while a check is made to see if a key has been pressed.

4.1.4 Event Services

4.1.4.1 gBS->CreateEvent(), gBS->SetTimer(), gBS->SignalEvent(), and gBS->CloseEvent()

The **gBS->CreateEvent()** and **gBS->CloseEvent()** services are used to create and destroy events. The following two basic types of events can be created:

- **EVT_NOTIFY_SIGNAL**
- **EVT_NOTIFY_WAIT**

The type of event determines when an event's notification function is invoked. The notification function for signal type events are invoked when the event is placed into the signaled state with a call to **gBS->SignalEvent()**. The notification function for wait type events are invoked when the event is passed to the **gBS->CheckEvent()** or **gBS->WaitForEvent()** services.

Some EFI drivers need to place their controllers in a quiescent state or perform other controller-specific actions at the time that an operating system is about to take full control of the platform. In this case, the EFI driver should create a signal type event that is notified when **gBS->ExitBootServices()** is called by the operating system. The notification function for this type of event is not allowed to use any of the EFI Memory Services either directly or indirectly because using those services may modify the memory map, which will force an error to be returned from **gBS->ExitBootServices()**. The code fragment in Example 4-17 shows how an Exit Boot Services event can be created and destroyed along with the skeleton of the notification function that is invoked when the Exit Boot Services event is signaled. This event is automatically signaled by EFI firmware when **gBS->ExitBootServices()** is called by an OS loader or an OS kernel. The notification function that is registered in **gBS->CreateEvent()** is called when the event is signaled.

```
VOID
EFIAPI
NotifyExitBootServices (
  IN EFI_EVENT  Event,
  IN VOID       *Context
  )
{
  //
  // Put driver-specific actions here.
  // No EFI Memory Service may be used directly or indirectly.
  //
}

EFI_STATUS  Status;
EFI_EVENT   *ExitBootServicesEvent;

//
// Create an Exit Boot Services event.
//
Status = gBS->CreateEvent (
                EFI_EVENT_SIGNAL_EXIT_BOOT_SERVICES,
                EFI_TPL_NOTIFY,
                NotifyExitBootServices,
                NULL,
                &ExitBootServicesEvent
                );
```

```
if (EFI_ERROR (Status)) {
  return Status;
}

//
// Destroy the Exit Boot Services event
//
Status = gBS->CloseEvent (ExitBootServicesEvent);
if (EFI_ERROR (Status)) {
  return Status;
}
```

Example 4-17. Exit Boot Services Event

EFI runtime drivers may need to be notified when the operating system sets up virtual mappings for all the EFI runtime mapping. In this case, the EFI driver should create a signal type event that is notified when **gBS->SetVirtualAddressMap()** is called by the operating system. This call allows the EFI runtime driver to convert pointers from physical addresses to virtual addresses. The notification function for this type of event is not allowed to use any of the EFI Boot Services, EFI Console Services, or EFI Protocol Services either directly or indirectly because those services are no longer available when **gRT->SetVirtualAddressMap()** is called. Instead, this type of notification function typically uses **gRT->ConvertPointer()** to convert pointers within data structures that are managed by the EFI runtime driver from physical addresses to virtual addresses. The code fragment in Example 4-18 shows how a Set Virtual Address Map event can be created and destroyed along with the skeleton of the notification function that is invoked when the Set Virtual Address Map event is signaled. As an example, this notification function converts a single global pointer from a physical address to a virtual address. This event is automatically signaled by EFI firmware when **gRT->SetVirtualAddressMap()** is called by an OS loader or an OS kernel. The notification function that is registered in **gBS->CreateEvent()** is called when the event is signaled.

```
VOID  *gGlobalPointer;

VOID
EFIAPI
NotifySetVirtualAddressMap (
  IN EFI_EVENT  Event,
  IN VOID       *Context
  )
{
  gRT->ConvertPointer (
         EFI_OPTIONAL_POINTER,
         (VOID **)&gGlobalPointer
         );
}

EFI_STATUS  Status;
EFI_EVENT   *SetVirtualAddressMapEvent;

//
// Create a Set Virtual Address Map event.
//
Status = gBS->CreateEvent (
                EFI_EVENT_SIGNAL_SET_VIRTUAL_ADDRESS_MAP,
                EFI_TPL_NOTIFY,
                NotifySetVirtualAddressMap,
```

```
                 NULL,
                 &SetVirtualAddressMapEvent
                 );
if (EFI_ERROR (Status)) {
  return Status;
}

//
// Destroy the Set Virtual Address Map event
//
Status = gBS->CloseEvent (SetVirtualAddressMapEvent);
if (EFI_ERROR (Status)) {
  return Status;
}
```

Example 4-18. Set Virtual Address Map Event

Some EFI drivers need to be notified on a periodic basis to poll a device. A good example is the USB bus driver, which needs to periodically check the status of all the ports to see if a USB device has been added or removed. Other drivers may need to be notified once after a specific period of time. A good example of this second scenario is the ISA floppy driver than needs to turn off the floppy drive motor if there are no read or write operations for a few seconds. These EFI drivers would use the **gBS->CreateEvent()** and **gBS->CloseEvent()** services to create and destroy the event and the **gBS->SetTimer()** to arm a periodic timer or a one-shot timer. See section 4.4.2 for examples on how to create periodic timer events and one-shot timer events.

Wait events are the last type of event that will be discussed here. Wait events are typically produced by protocols that abstract I/O services for controllers that provide input services. These protocols will contain an **EFI_EVENT** data field that is signaled by the EFI driver when the input is available. The following protocols in the *EFI 1.10 Specification* use this mechanism:

- **EFI_SIMPLE_INPUT_PROTOCOL** (section 10.2 in the *EFI 1.10 Specification*)
- **EFI_SIMPLE_POINTER_PROTOCOL** (section 10.10 in the *EFI 1.10 Specification*)
- **EFI_SIMPLE_NETWORK_PROTOCOL** (section 15.1 in the *EFI 1.10 Specification*)

The code fragment in Example 4-19 shows an example of a wait event that is created by a keyboard driver that produces the **EFI_SIMPLE_INPUT_PROTOCOL**. The first part of the code fragment is the event notification function that is called when the wait event is waited on with the **gBS->CheckEvent()** or the **gBS->WaitForEvent()** services. The second part of the code fragment is the code that would be in the **Start()** and **Stop()** functions that create and destroy the wait event. Typically, an EFI application or the EFI boot manager will call **gBS->CheckEvent()** or **gBS->WaitForEvent()** to see if a key has been pressed on a input device that supports the **EFI_SIMPLE_INPUT_PROTOCOL**. This call to **gBS->CheckEvent()** or **gBS->WaitForEvcent()** will cause the notification function of the wait event in the **EFI_SIMPLE_INPUT_PROTOCOL** to be executed. The notification function checks to see if a key has been pressed on the input device. If the key has been pressed, then the wait event is signaled with a call to **gBS->SignalEvent()**. If the wait event is signaled, then the EFI application or EFI boot manager will get an **EFI_SUCCESS** return code, and the EFI application or EFI boot manager can call the **ReadKeyStroke()** service of the **EFI_SIMPLE_INPUT_PROTOCOL** to read the key that was pressed.

```
VOID
EFIAPI
```

```
KeyboardWaitForKey (
  IN  EFI_EVENT  Event,
  IN  VOID       *Context
  )
{
  EFI_TPL  OldTpl;

  //
  // Raise the Task Priority Level to EFI_TPL_NOTIFY to perform blocking I/O
  //
  OldTpl = gBS->RaiseTPL (EFI_TPL_NOTIFY);

  //
  // Call an internal function to see if a key has been pressed
  //
  if (!EFI_ERROR (KeyboardCheckForKey (Context))) {
    //
    // If a key has been pressed, then signal the wait event
    //
    gBS->SignalEvent (Event);
  }

  //
  // Restore the Task Priority Level to its original level
  //
  gBS -> RestoreTPL (OldTpl);

  return;
}

EFI_STATUS                  Status;
EFI_SIMPLE_INPUT_PROTOCOL  *SimpleInput;

//
// Create a wait event for a Simple Input Protocol
//
Status = gBS->CreateEvent (
                EFI_EVENT_NOTIFY_WAIT,
                EFI_TPL_NOTIFY,
                KeyboardWaitForKey,
                &SimpleInput,
                &SimpleInput.WaitForKey
                );
if (EFI_ERROR (Status)) {
  return Status;
}

//
// Destroy the wait event
//
Status = gBS->CloseEvent (SimpleInput.WaitForKey);
if (EFI_ERROR (Status)) {
  return Status;
}
```

Example 4-19. Wait for Event

4.1.4.2 gBS->CheckEvent()

This service checks to see if an event is in the waiting state or the signaled state. EFI drivers will typically use this service to see if a one-shot timer event or a wait event is in the signaled state. For example, the PXE base code driver and the terminal driver create a one-shot timer event and they both use this service to see if a transaction has timed out. The console splitter driver uses this service to see if a key has been pressed or a pointer device has been moved on one of the input devices that it is managing. The code fragment in Example 4-20 creates a one-shot timer event and uses **gBS->CheckEvent()** to wait until the timer expires.

```
EFI_STATUS  Status;
EFI_EVENT   TimerEvent;

Status = gBS->CreateEvent (
                EFI_EVENT_TIMER | EFI_EVENT_NOTIFY_WAIT,
                EFI_TPL_NOTIFY,
                NULL,
                NULL,
                &TimerEvent
                );

Status = gBS->SetTimer (
                TimerEvent,
                TimerRelative,
                40000000         // 4 seconds in the future
                );

do {
  Status = gBS->CheckEvent (TimerEvent);
} while (EFI_ERROR (Status));
```

Example 4-20. Wait for a One-Shot Timer Event

4.1.5 Delay Services

4.1.5.1 gBS->Stall()

This service waits for a specified number of microseconds. The range of supported delays is from 1 µS to 4294 seconds. However, the delays passed into this service should be short and are typically in the range of a few microseconds to a few milliseconds. See section 4.4.2 for examples of how this service can be used.

4.2 Services That EFI Drivers Rarely Use

Table 4-2 contains the list of EFI services that are rarely used by EFI drivers. The following sections provide a brief description of each service, why they are rarely used, and a code example on how they are typically used by EFI drivers. This table is also present in Appendix C for reference purposes.

Table 4-2. EFI Services That Are Rarely Used by EFI Drivers

Type	Service	Type	Service
BS	gBS->ReinstallProtocolInterface()	BS	gBS->LoadImage()
BS	gBS->LocateDevicePath()	BS	gBS->StartImage()
		BS	gBS->UnloadImage()
BS	gBS->ConnectController()	BS	gBS->Exit()
BS	gBS->DisconnectController()		
		BS	gBS->InstallConfigurationTable()
RT	gRT->GetVariable()		
RT	gRT->SetVariable()	RT	gRT->GetTime()
BS	gBS->GetNextMonotonicCount()	BS	gBS->CalculateCrc32()
RT	gRT->GetNextHighMonotonicCount()		
		RT	gRT->ConvertPointer()

4.2.1 Handle Database and Protocol Services

4.2.1.1 gBS->ReinstallProtocolInterface()

This service should be used only to indicate media change events and when a device path is modified or updated. This service applies to the following:

- The Block I/O Protocol when the media in a removable media device is changed
- The Serial I/O Protocol when its attributes are modified with a call to **`SetAttributes()`**
- The Simple Network Protocol when the MAC address of the network interface is modified with a call to **`StationAddress()`**

This service is basically a series of the following calls, in the order listed:

1. **`UninstallProtocolInterface()`**, which may cause **`DisconnectController()`** to be called
2. **`InstallProtocolInterface()`**
3. **`ConnectConroller()`** to allow controllers that had to release the protocol a chance to connect to it again

The code fragment in Example 4-21 shows what an EFI driver that produces the Block I/O Protocol would do when the media in a removable media device is changed. The exact same protocol is reinstalled onto the controller handle.

```
EFI_STATUS              Status;
EFI_HANDLE              ControllerHandle;
EFI_BLOCK_IO_PROTOCOL   *BlockIo;

Status = gBS->ReinstallProtocolInterface (
                ControllerHandle,
                &gEfiBlockIoProtocolGuid,
                BlockIo,
                BlockIo
                );
```

Example 4-21. Reinstall Protocol Interface

NOTE

This service can cause reentrancy problems if not handled correctly. If a driver makes a request that requires a protocol of a parent device to be updated, then that protocol will be removed and reattached. The driver making the request may not realize that the request will cause the driver to be completely stopped and completely restarted when the request to the parent device is made. For example, consider a terminal driver that wants to change the baud rate on the serial port. The baud rate is changed with a call to the Serial I/O Protocol's **SetAttribute()**. *This call changes the baud rate, which is reflected in the device path of the serial device, so the Device Path Protocol is reinstalled by the* **SetAttributes()** *service. This reinstallation will force the terminal driver to be disconnected. The terminal driver will then attempt to connect to the serial device again, but the baud rate will be the one that the terminal driver expects, so the terminal driver will not need to set the baud rate again. Any consumer of a protocol that supports this media change concept needs to be aware that the protocol can be reinstalled at any time and care must be taken in the design of drivers that use this type of protocol.*

4.2.1.2 gBS->LocateDevicePath()

This service locates a device handle that supports a specific protocol and has the closest matching device path. This service is useful when an EFI driver needs to find an I/O abstraction for one of its parent controllers. Normally, an EFI driver will use the services on the *ControllerHandle* that is passed into the **Supported()** and **Start()** functions of the EFI driver's **EFI_DRIVER_BINDING_PROTOCOL**. However, if an EFI driver needs to use services from a parent controller, this function can be used to find the handle of a parent controller. For example, a PCI device driver will normally use the PCI I/O Protocol to manage a PCI controller. If the PCI device driver needs the services of the PCI Root Bridge I/O Protocol of which the PCI controller is a child, then the **gBS->LocateDevicePath()** function can be used to find the parent handle that supports the PCI Root Bridge I/O Protocol, and then the **gBS->OpenProtocol()** service can be used to retrieve the PCI Root Bridge I/O Protocol interface. This operation is not recommended because a parent bus driver typically owns the parent I/O abstractions. Directly using a parent I/O may cause unintended side effects. The code fragment in Example 4-22 demonstrates this example.

Section 14.4.2 contains another example that shows the recommended method for a PCI driver to access the resources of different PCI controllers without using the PCI Root Bridge I/O Protocol.

```
EFI_STATUS                        Status;
EFI_GUID                          gEfiDevicePathProtocolGuid;
EFI_GUID                          gEfiPciRootBridgeIoProtocolGuid;
EFI_DRIVER_BINDING_PROTOCOL       *This;
EFI_HANDLE                        ControllerHandle;
EFI_DEVICE_PATH_PROTOCOL          *DevicePath;
EFI_HANDLE                        ParentHandle;
EFI_PCI_ROOT_BRIDGE_IO_PROTOCOL   *PciRootBridgeIo;

//
// Retrieve the Device Path Protocol instance on ControllerHandle
//
Status = gBS->OpenProtocol (
                ControllerHandle,
                &gEfiDevicePathProtocolGuid,
                &DevicePath,
                This->DriverBindingHandle,
                ControllerHandle,
                EFI_OPEN_PROTOCOL_GET_PROTOCOL
                );
if (EFI_ERROR (Status)) {
  return Status;
}


//
// Find a parent controller that supports the PCI Root Bridge I/O Protocol
//
Status = gBS->LocateDevicePath (
                 &gEfiPciRootBridgeIoProtocolGuid,
                 &DevicePath,
                 &ParentHandle
                 );
if (EFI_ERROR (Status)) {
  return Status;
}

//
// Get the PCI Root Bridge I/O Protocol instance on ParentHandle
//
Status = gBS->OpenProtocol (
                 ParentHandle,
                 &gEfiPciRootBridgeIoProtocolGuid,
                 &PciRootBridgeIo,
                 This->DriverBindingHandle,
                 ControllerHandle,
                 EFI_OPEN_PROTOCOL_GET_PROTOCOL
                 );
if (EFI_ERROR (Status)) {
  return Status;
}
```

Example 4-22. Locate Device Path

4.2.1.3 gBS->ConnectController() and gBS->DisconnectController()

These services are used by EFI drivers that are bus drivers for bus types with hot-plug capabilities that are supported in the preboot environment. The only bus driver in the *EFI Sample Implementation* that uses these services is the USB bus driver. The USB bus driver does not create any child handles in its **Start()** function. Instead, it registers a periodic timer event. Each time the timer period expires, the timer event's notification function is called, and that notification function examines all the USB hubs to see if any USB devices have been added or removed. If a USB device has been added, then a child handle is created, and **gBS->ConnectController()** is called so the USB device drivers can connect to the newly added USB device. If a USB device has been removed, then **gBS->DisconnectController()** is called to stop the USB device drivers from managing the USB device that was just removed. Just because a bus is capable of supporting hot-plug events does not necessarily mean that the EFI driver for that bus type must support those hot-plug events. Support for hot-plug events in the preboot environment is dependent on the platform requirements for each bus type. The code fragment in Example 4-23 shows how these services could be used from the USB bus driver.

```
EFI_STATUS  Status;
BOOLEAN     HotAdd;
BOOLEAN     HotRemove;
EFI_HANDLE  ChildHandle;

//
// If ChildHandle is a device that was just hot added, then recursively
// connect all drivers to ChildHandle
//
if (HotAdd == TRUE) {
  Status = gBS->ConnectController(
                  ChildHandle,
                  NULL,
                  NULL,
                  TRUE
                  );
}

//
// If ChildHandle is a device that was just removed, then recursively
// disconnect all drivers from ChildHandle
//
if (HotRemove == TRUE) {
  Status = gBS->DisconnectController(
                  ChildHandle,
                  NULL,
                  NULL
                  );
}
```

Example 4-23. Connect and Disconnect Controller

The **gBS->DisconnectController()** service may also be used from unloadable EFI drivers to disconnect the EFI driver from the device it is managing in its **Unload()** function. The code fragment in Example 4-24 shows one algorithm that an **Unload()** function can use to disconnect the driver from all the devices in the system. It retrieves the list of all the handles in the handle database and disconnects the EFI driver from each of those handles and frees the buffer containing the list handles.

```
EFI_STATUS  Status;
EFI_HANDLE  ImageHandle;
EFI_HANDLE  *DeviceHandleBuffer;
UINTN       DeviceHandleCount;
UINTN       Index;

Status = gBS->LocateHandleBuffer (
                AllHandles,
                NULL,
                NULL,
                &DeviceHandleCount,
                &DeviceHandleBuffer
                );
if (EFI_ERROR (Status)) {
  return Status;
}

for (Index = 0; Index < DeviceHandleCount; Index++) {
  Status = gBS->DisconnectController (
                  DeviceHandleBuffer[Index],
                  ImageHandle,
                  NULL
                  );
}

gBS->FreePool(DeviceHandleBuffer);
```

Example 4-24. Disconnect Controller

4.2.2 Image Services

4.2.2.1 gBS->LoadImage(), gBS->StartImage(), gBS->UnloadImage(), and gBS->Exit()

EFI drivers should not call **gBS->Exit()**. Instead, they should just return a status code from their driver entry point.

EFI drivers do not normally load and unload other EFI drivers or EFI applications. However, there are two exceptions. The first exception is that bus drivers that can load, start, and potentially unload EFI drivers that are stored in containers on the child devices of the bus. For example, the PCI bus driver loads and starts EFI drivers that are stored in PCI option ROMs.

The second exception is for devices that have an EFI driver that manages the device and an EFI application that provides the user interface that is used to configure the device. In this case, the implementation of the **SetOptions()** service in the **EFI_DRIVER_CONFIGURATION_PROTOCOL** uses the **gBS->LoadImage()** and **gBS->StartImage()** services to load and execute the EFI application to configure the device.

4.2.3 Variable Services

4.2.3.1 gRT->GetVariable() and gRT->SetVariable()

These services are used to get and set EFI variables. When variables are stored, there are attributes that describe the visibility and persistence of the variable. These attributes can be combined different types of variables. The legal combinations of attributes include the following:

BOOTSERVICE_ACCESS The variable is available for read and write access in the preboot environment before **gBS->ExitBootServices()** is called. The variable is not available after **gBS->ExitBootServices()** is called, and contents are also lost on the next system reset or power cycle. These types of variables are typically used to share information among different preboot components.

BOOTSERVICE_ACCESS | RUNTIME_ACCESS
The variable is available for read and write access in the preboot environment before **gBS->ExitBootServices()** is called. It is available for read-only access from the OS runtime environment after **gBS->ExitBootServices()** is called. The contents are lost on the next system reset or power cycle. These types of variable are typically used to share information among different preboot components and pass read-only information to the operating system.

NON_VOLATILE | BOOTSERVICE_ACCESS
The variable is available for read and write access in the preboot environment before **gBS->ExitBootServices()** is called, and the contents are persistent across system resets and power cycles. These types of variables are typically used to share persistent information among different preboot components.

NON_VOLATILE | BOOTSERVICE_ACCESS | RUNTIME_ACCESS
The variable is available for read and write access in both the preboot environment and the OS runtime environment. The contents are persistent across system resets and power cycles. These types of variables are typically used to share persistent information among preboot components and the operating system.

The code fragment in Example 4-25 shows how a fixed-sized EFI variable can be read and written. A variable name is a combination of a GUID and a Unicode string. The GUID allows private variables to be managed by different vendors. Section 3.2 of the *EFI 1.10 Specification* defines one GUID that is used to access EFI variables. This example shows how the fixed-size EFI variable *BootNext* can be accessed.

```
EFI_STATUS  Status;
UINT32      Attrubutes;
UINTN       DataSize;
UINT16      BootNext;

DataSize = sizeof (UINT16);
Status = gRT->GetVariable (
                L"BootNext",
                &gEfiGlobalVariableGuid,
                &Attributes,
                &DataSize,
                &BootNext
                );

BootNext = 2;
Status = gRT->SetVariable (
                L"BootNext",
                &gEfiGlobalVariableGuid,
                EFI_VARIABLE_NON_VOLATILE |
                  EFI_VARIABLE_BOOTSERVICE_ACCESS |
                  EFI_VARIABLE_RUNTIME_ACCESS,
                sizeof (UINT16),
                &BootNext
                );
```

Example 4-25. Reading and Writing Fixed-Size EFI Variables

The code fragment in Example 4-26 shows how a memory buffer for the variable-sized EFI variable *Driver001C* can be allocated. The variable is then read into the allocated buffer and written to the EFI variable *Driver001D*. Finally, the allocated buffer is freed.

```
EFI_STATUS  Status;
UINT32      Attrubutes;
UINTN       DataSize;
VOID        *Data;

DataSize = 0;
Status = gRT->GetVariable (
                L"Driver001C",
                &gEfiGlobalVariableGuid,
                &Attributes,
                &DataSize,
                Data
                );
if (Status != EFI_BUFFER_TOO_SMALL) {
  return Status;
}

Data = EfiLibAllocatePool (DataSize);
if (Data == NULL) {
  return EFI_OUT_OF_RESOURCES;
}

Status = gRT->GetVariable (
                L"Driver001C",
                &gEfiGlobalVariableGuid,
                &Attributes,
                &DataSize,
                Data
```

```
                 );
if (EFI_ERROR (Status)) {
  return Status;
}

Status = gRT->SetVariable (
                L"Driver001D",
                &gEfiGlobalVariableGuid,
                EFI_VARIABLE_NON_VOLATILE |
                  EFI_VARIABLE_BOOTSERVICE_ACCESS |
                  EFI_VARIABLE_RUNTIME_ACCESS,
                DataSize,
                Data
                );

gBS->FreePool (Data);
```

Example 4-26. Reading and Writing Variable-Sized EFI Variables

4.2.4 Time Services

4.2.4.1 gRT->GetTime()

This service is typically only accurate to about 1 second. As a result, EFI drivers should not use this service to poll or wait for an event from a device. Instead, the **gBS->Stall()** service and the **gBS->SetTimer()** services provide time services with much higher accuracy. This service should be used only when the current time and date are required such as recording the time and date of a critical error. See section 4.4.3 for more details on the time-related services that are provided by EFI.

4.2.5 Virtual Memory Services

4.2.5.1 gRT->ConvertPointer()

This service is never used by EFI Boot Service drivers and is sometimes used by EFI runtime drivers. Example 4-18 shows an example of how an EFI runtime driver might use this service.

4.2.6 Miscellaneous Services

4.2.6.1 gBS->InstallConfigurationTable()

This service is used to add, update, or remove an entry in the list of configuration tables that is maintained in the EFI System Table. These configuration tables are typically used to pass information from the EFI environment into an operating system environment. EFI Boot Service drivers are destroyed at **gBS->ExitBootServices()**, so they do not typically need to pass any information to an operating system. EFI runtime drivers continue to persist after **gBS->ExitBootServices()**, so they may need to pass information to an operating system so that the operating system can use the services that the EFI runtime driver produced. Typically, only the class of EFI runtime drivers that need to pass information to the operating system will use this service. The code fragment in Example 4-27 shows how an UNDI driver can add, update, and remove a configuration table. The first parameter is a pointer to the Network Interface Identifier

(NII) GUID, and the second parameter is a pointer to the data structure that is associated with the NII GUID.

```
EFI_STATUS  Status;
NII_TABLE   *UndiData;

//
// Add or update a configuration table
//
Status = gBS->InstallConfigurationTable (
                &gEfiNetworkInterfaceIdentifierProtocolGuid_31,
                UndiData
                );

//
// Remove a configuration table
//
Status = gBS->InstallConfigurationTable (
                &gEfiNetworkInterfaceIdentifierProtocolGuid_31,
                NULL
                );
```

Example 4-27. Install Configuration Table

4.2.6.2 gBS->CalculateCrc32()

This service is used by the EFI core to maintain the checksum of the EFI System Table, EFI Boot Services Table, and EFI Runtime Services Table. It is also used by a few EFI drivers including the Console Splitter and the Partition driver for Guided Partition Table (GPT) disk. If an EFI driver requires the use of 32-bit CRCs, the size of the EFI driver can be reduced by using this EFI boot service. The code fragment in Example 4-28 shows how this service can be used to calculate a 32-bit CRC value for *Size* bytes of a buffer that has a header of type **EFI_TABLE_HEADER**. **EFI_TABLE_HEADER** is defined in section 4.2 of the *EFI 1.10 Specification*, and it is the standard header used for the EFI System Table, EFI Boot Services Table, EFI Runtime Services Table, and GPT related structures.

```
EFI_STATUS        Status;
UINT32            Crc;
EFI_TABLE_HEADER  *Header;
UINT32            Size;

Header->CRC32 = 0;
Status = gBS->CalculateCrc32(
                (VOID *)Header,
                Size,
                &Crc
                );
Hdr->CRC32 = Crc;
```

Example 4-28. Computing 32-bit CRC Values

This service can also be used to verify a 32-bit CRC value. The code fragment in Example 4-29 shows how the 32-bit CRC for a buffer of *Size* bytes with an **EFI_TABLE_HEADER** can be validated. It returns **TRUE** if the 32-bit CRC is good. Otherwise, it returns **FALSE**.

```
EFI_STATUS        Status;
UINT32            OriginalCrc;
```

```
UINT32            Crc;
EFI_TABLE_HEADER  *Header;
UINT32            Size;

OriginalCrc = Header->CRC32;
Header->CRC32 = 0;
Status = gBS->CalculateCrc32(
                (VOID *)Header,
                Size,
                &Crc
                );
Hdr->CRC32 = OriginalCrc;

if (OriginalCrc == Crc) {
  return TRUE;
} else {
  return FALSE;
}
```

Example 4-29. Validating 32-bit CRC Values

4.2.6.3 gBS->GetNextMonotonicCount() and gRT->GetNextHighMonotonicCount()

These services provide a 64-bit monotonic counter that is guaranteed to increase. They are not used by any of the drivers in the *EFI Sample Implementation*.

4.3 Services That EFI Drivers Should Not Use

Table 4-3 lists the EFI services that should not be used by EFI drivers. The following sections describe why each of these functions should not be used. This table is also present in Appendix C for reference purposes.

Table 4-3. EFI Services That Should Not Be Used by EFI Drivers

Type	Service	Type	Service
BS	gBS->GetMemoryMap()	RT	gRT->SetVirtualAddressMap()
BS	gBS->ExitBootServices()	RT	gRT->GetNextVariableName()
BS	gBS->InstallProtocolInterface()	RT	gRT->SetTime()
BS	gBS->UninstallProtocolInterface()	RT	gRT->GetWakeupTime()
BS	gBS->HandleProtocol()	RT	gRT->SetWakeupTime()
BS	gBS->LocateHandle()		
BS	gBS->RegisterProtocolNotify()	RT	gRT->ResetSystem()
BS	gBS->ProtocolsPerHandle()	BS	gBS->SetWatchDogTimer()
BS	gBS->WaitForEvent()		

4.3.1 Memory Services

4.3.1.1 gBS->GetMemoryMap()

EFI drivers should not use this service because EFI drivers should not depend upon the physical memory map of the platform. The **gBS->AllocatePool()** and **gBS->AllocatePages()** services allow an EFI driver to allocate system memory. The **gBS->FreePool()** and **gBS->FreePages()** allow an EFI driver to free previously allocated memory. If there are limitations on the memory areas that a specific device may use, then those limitations should be managed by a parent I/O abstraction that understands the details of the platform hardware. For example, PCI device drivers should use the services of the PCI I/O Protocol to manage DMA buffers. The PCI I/O Protocol is produced by the PCI bus driver that uses the services if the PCI Root Bridge I/O Protocol to manage DMA buffers. The PCI Root Bridge I/O Protocol is chipset and platform specific, so the component that produces the PCI Root Bridge I/O Protocol understands what memory regions can be used for DMA operations. By pushing the responsibility into the chipset- and platform-specific components, the PCI device drivers and PCI bus drivers are easier to implement and are portable across a wide variety of platforms.

4.3.2 Image Services

4.3.2.1 gBS->ExitBootServices()

This service is used only by EFI OS loaders or OS kernels. It hands control of the platform from the EFI to an OS. After this call, the EFI Boot Services are no longer available, and all memory in use by EFI Boot Service drivers is considered to be available memory. If this call is made by an EFI Boot Service driver, it would essentially destroy itself.

4.3.3 Handle Database and Protocol Services

4.3.3.1 gBS->InstallProtocolInterface()

This service adds one protocol interface to an existing handle or creates a new handle. This service has been replaced by the **gBS->InstallMultipleProtocolInterfaces()** service, so all EFI drivers should use the newer service. Using this replacement service provides additional flexibility and additional error checking and produces smaller EFI drivers.

4.3.3.2 gBS->UninstallProtocolInterface()

This service removes one protocol interface from a handle in the handle database. The functionality of this service has been replaced by **gBS->UninstallMultipleProtocolInterfaces()**. This service uninstalls one or more protocol interfaces from the same handle. Using this replacement service provides additional flexibility and produces smaller EFI drivers.

4.3.3.3 gBS->HandleProtocol()

EFI drivers must not use this service because they will not be compliant with the EFI Driver Model if they do, which could introduce interoperability issues. Instead, **`gBS->OpenProtocol()`** should be used because it provides the equivalent functionality, and it allows the EFI core to track the agents that are using different protocol interfaces in the handle database.

4.3.3.4 gBS->LocateHandle()

This service returns an array of handles that support the specified protocol. This service requires the caller to allocate the return buffer. The **`gBS->LocateHandleBuffer()`** service is easier to use and produces smaller executables because it allocates the return buffer for the caller.

4.3.3.5 gBS->RegisterProtocolNotify()

This service registers an event that is to be signaled whenever an interface is installed for a specified protocol. Using this service is strongly discouraged. This service was previously used by EFI drivers that follow the *EFI 1.02 Specification*, and it provided a simple mechanism for drivers to layer on top of another driver. Chapter 9 of the *EFI 1.10 Specification* instead defines the EFI Driver Model, which provides a much more flexible mechanism.

4.3.3.6 gBS->ProtocolsPerHandle()

This service retrieves the list of protocols that are installed on a handle. Because EFI drivers should already know what protocols are installed on the handles that the EFI driver is managing, this service should not be used. This service is typically used by EFI applications that need to traverse the entire handle database.

4.3.4 Event Services

4.3.4.1 gBS->WaitForEvent()

This service waits for an event in an event list to be signaled. EFI drivers are typically waiting only for a single event to enter the signaled state, so the **`gBS->CheckEvent()`** service should be used instead.

4.3.5 Virtual Memory Services

4.3.5.1 gBS->SetVirtualAddressMap()

This service is also used only by EFI OS loaders or OS kernels for operating systems that wish to call EFI Runtime Services using virtual addresses. This service must be called after **`gBS->ExitBootServices()`** is called. As a result, it is not legal for EFI drivers to call this service.

4.3.6 Variables Services

4.3.6.1 gRT->GetNextVariableName()

This service is used to walk the list of EFI variables that are maintained through the EFI Variable Services. Most EFI drivers should already know the EFI variables they want to access, so there is no need for an EFI driver to walk the list of all the EFI variables.

4.3.7 Time Services

4.3.7.1 gRT->SetTime(), gRT->GetWakeupTime(), and gRT->SetWakeupTime()

EFI drivers should not modify the system time or the wakeup timer. The management of these timer services should be left to the EFI boot manager, an OEM-provided utility, or an operating system.

4.3.8 Miscellaneous Services

4.3.8.1 gBS->SetWatchdogTimer()

The watchdog timer is managed from the EFI boot manager, so EFI drivers should not use this service.

4.3.8.2 gRT->ResetSystem()

System resets should be managed from the EFI boot manager or OEM-provided utilities. EFI drivers should not use this service. The only exceptions in the *EFI Sample Implementation* are the keyboard drivers that detect the CTRL-ALT-DEL key sequence to reset the platform.

4.4 Time-Related Services

There are several different time-related services that are available to EFI drivers, and they are listed below in Table 4-4. The time-related services that should not be used by EFI drivers are discussed in section 4.4.2 and include the following:

- `gRT->SetTime()`
- `gRT->GetWakeupTime()`
- `gRT->SetWakeupTime()`
- `gBS->SetWatchDogTimer()`

Because EFI drivers should not use these services, they will not be discussed. Omitting the four services above leaves the following three time-related services that EFI drivers can use:

- `gBS->SetTimer()`
- `gBS->Stall()`
- `gRT->GetTime()`

The **gBS->SetTimer()** and **gBS->Stall()** services are commonly used, but the **gRT->GetTime()** service is rarely used.

If an EFI driver requires highly accurate short delays, then the **gBS->Stall()** service should be used. If an EFI driver needs to periodically synchronize with a device, then the **gBS->SetTimer()** service should be used with an event. If an EFI driver needs to know the current time and date, then the **gRT->GetTime()** service should be used.

Table 4-4. Time-Related EFI Services

Type	Service	Type	Service
BS	gBS->SetTimer()	RT	gRT->GetWakeupTime()
BS	gBS->Stall()	RT	gRT->SetWakeupTime()
RT	gRT->GetTime()	BS	gBS->SetWatchDogTimer()
RT	gRT->SetTime()		

4.4.1 Stall Service

The **gBS->Stall()** service is the time-related service with the highest accuracy. The stall times can range from 1 µS to about 4294 seconds. Most implementations of the stall service use a calibrated software loop, so they are very accurate. This service is a good one to use when an EFI driver requires a short delay. For example, an EFI driver may send a command to a controller and then wait for the command to complete. Because all EFI drivers are polled, the EFI driver could use the stall service inside a loop to periodically check for the completion status (see Example 4-30). Another good use of the stall service is for hardware devices that require delays between register accesses. Here, a fixed stall value would be used, and the stall value would be based in a hardware specification for the device that is being accessed (see Example 4-31). One disadvantage of the stall service is that other EFI components cannot execute while a stall is being executed. If long delays are required and it makes sense to defer the completion of an I/O operation, then the timer event services described in the next section should be used.

```
EFI_STATUS            Status;
UINTN                 TimeOut;
EFI_PCI_IO_PROTOCOL   PciIo;
UITN8                 Value;

//
// Loop waiting for the register at Offset 0 of Bar #0 of PciIo to become 0xE0.
// Wait 10 uS between each check of this register, and time out if it does
// not become 0 after 100 mS.
//
TimeOut = 0;
do {
  //
  // Wait 10 uS
  //
  gBS->Stall(10);

  //
  // Increment TimeOut by the number of stalled uS
  //
  TimeOut = TimeOut + 10;
```

```
  //
  // Do a single 8 bit read from BAR #0, Offset 0 into Value
  //
  Status = PciIo->Io.Read (
                       PciIo,                // This
                       EfiPciIoWidthUint8,   // Width
                       0,                    // BarIndex
                       0,                    // Offset
                       1,                    // Count
                       &Value                // Buffer
                       );
} while (Value != 0xE0 && TimeOut <= 100000);

if (Value != 0xE0) {
  return EFI_TIMEOUT;
}

return EFI_SUCCESS;
```

Example 4-30. Stall Loop

```
EFI_STATUS           Status;
EFI_PCI_IO_PROTOCOL  PciIo;
UITN8                Value;

//
// This example shows a stall of 1000 uS between two register writes to the
// same register. The 1000 uS is based on a hardware requirement for the
// device being accessed.
//

//
// Do a single 8 bit write to BAR #1, Offset 0x10 of 0xAA
//
Value = 0xAA;
Status = PciIo->Io.Write (
                     PciIo,                // This
                     EfiPciIoWidthUint8,   // Width
                     1,                    // BarIndex
                     0x10,                 // Offset
                     1,                    // Count
                     &Value                // Buffer
                     );

//
// Wait 1000 uS
//
gBS->Stall(1000);

//
// Do a single 8-bit write to BAR #1, Offset 0x10 of 0x55
//
Value = 0x55;
Status = PciIo->Io.Write (
                     PciIo,                // This
                     EfiPciIoWidthUint8,   // Width
                     1,                    // BarIndex
```

```
                    0x10,                   // Offset
                    1,                      // Count
                    &Value                  // Buffer
                    );
```

Example 4-31. Stall Service

4.4.2 Timer Events

The timer event services allow an EFI driver to wait for a specified period of time before a notification function is called. This service is a good choice for EFI drivers that need to defer the completion of an I/O operation, thus allowing other EFI components to continue to execute while the EFI driver waits for the specified period of time. The time is specified in 100 nS units. This unit may give the appearance of having better accuracy than the **gBS->Stall()** service, which has an accuracy of 1 µS, but that is not the case. The EFI core uses a single timer interrupt to determine when to signal timer events. The resolution of timer events is completely dependent on the frequency of the timer interrupt that is used by the EFI core. In most systems, the timer interrupt is generated every 10 to 50 mS, but the *EFI 1.10 Specification* does not require any specific interrupt rate. This lack of specificity means that a periodic timer that is set with a 100 nS period will actually get called only every 10 mS to 50 mS, which is why the **gBS->Stall()** service is a much better choice for short delays.

However, there are cases when timer events work very well. One example is the USB bus driver. This driver is required to periodically check the status of all the USB hubs to see if any USB devices have been attached or removed. This operation does not need to be performed very frequently, so the timer events work very well. Example 4-32 shows how to set up a periodic timer event with a period of 100 mS. This period is similar to what the USB bus driver would use. Another good example is a floppy driver. This driver needs to manage the drive motor on the floppy drive so that it is automatically turned off if there are no floppy transactions for a period of time. Example 4-33 shows how to set up a timer event that will fire 4 seconds in the future. If a timer like this is rearmed every time a floppy transaction is performed, then the notification function for this timer could be used to turn off the floppy drive motor.

```
//
// This is the notification function used in the peridic timer example below.
// An EFI driver will perform a driver-specific operation inside this function
// The Context parameter will typically contain a private context structure for
// a device that is managed by the EFI driver.
//
VOID
TimerHandler (
  IN EFI_EVENT  Event,
  IN VOID       *Context
  )

{
}

EFI_STATUS  Status;
EFI_EVENT   TimerEvent;
VOID        *Context;

//
// Create a timer event that will call the notification function TimerHandler()
```

```
// at a TPL level of EFI_TPL_NOTIFY when the timer fires. A context can be
// passed into the notification function when the timer fires, and that context
// is specified by the VOID pointer Context. When the event is created, it is
// returned in TimerEvent.
//
Status = gBS->CreateEvent (
                EFI_EVENT_TIMER | EFI_EVENT_NOTIFY_SIGNAL,
                EFI_TPL_NOTIFY,
                TimerHandler,
                Context,
                &TimerEvent
                );

//
// Set the timer value for the event created above to be a periodic timer with
// a period of 100 mS. The timer values are specified in 100 nS units.
// The notification function TimerHandler() will be called every 100 mS, and
// Context will be passed to the notification function.
//
Status = gBS->SetTimer (
                TimerEvent,
                TimerPeriodic,
                1000000          // Every 100 mS
                );
```

Example 4-32. Starting a Periodic Timer

```
//
// This is the notification function used in the one-shot timer example below.
// An EFI driver will perform a driver-specific operation inside this function.
// The Context parameter will typically contain a private context structure for
// a device that is managed by the EFI driver.
//
VOID
TimerHandler(
  IN EFI_EVENT  Event,
  IN VOID       *Context
  )

{
}

EFI_STATUS  Status;
EFI_EVENT   TimerEvent;
VOID        *Context;

//
// Create a timer event that will call the notification function TimerHandler()
// at a TPL level of EFI_TPL_NOTIFY when the timer fires. A context can be
// passed into the notification function when the timer fires, and that context
// is specified by the VOID pointer Context. When the event is created, it is
// returned in TimerEvent.
//
Status = gBS->CreateEvent (
                EFI_EVENT_TIMER | EFI_EVENT_NOTIFY_SIGNAL,
                EFI_TPL_NOTIFY,
                TimerHandler,
                Context,
                &TimerEvent
                );
```

```
//
// Set the timer value for the event created above to be a one-shot timer with
// a delay of 4 seconds. The timer values are specified in 100 nS units.
// The notification function TimerHandler() will be called 4 seconds after this
// call returns, and Context will be passed to the notification function.
//
Status = gBS->SetTimer (
                TimerEvent,
                TimerRelative,
                40000000          // 4 seconds in the future
                );
```

Example 4-33. Arming a One-Shot Timer

Once an EFI driver is finished with an event timer, the timer must be stopped. The **gBS->SetTimer()** service can be used to cancel the event timer, and the **gBS->CloseEvent()** service can be used to free the event that was allocated in **gBS->CreateEvent()**. If the timer was created in the **EFI_DRIVER_BINDING_PROTOCOL.Start()** function, then this operation would be performed in the **EFI_DRIVER_BINDING_PROTOCOL.Stop()** function. If the timer was created in the driver's entry point, then this operation would be performed in the driver's unload function if the driver has an unload function. Example 4-34 below shows how to cancel and free the timer event that was created for the examples in Example 4-32 and Example 4-33.

```
EFI_STATUS  Status;
EFI_EVENT   TimerEvent;

Status = gBS->SetTimer (
                TimerEvent,
                TimerCancel,
                0
                );

Status = gBS->CloseEvent (
                TimerEvent
                );
```

Example 4-34. Stopping a Timer

4.4.3 Time and Date Services

Example 4-35 shows two examples of the **gRT->GetTime()** service. The first retrieves the current time and date in an **EFI_TIME** structure and it retrieves the capabilities of the real-time clock hardware in an **EFI_TIME_CAPABILITIES** structure. The second example just retrieves the current time and date in an **EFI_TIME** structure and does not retrieve the capabilities of the real time clock hardware. See the *EFI 1.10 Specification* for a detailed description of this service and its associated data structures.

```
EFI_STATUS              Status;
EFI_TIME                Time;
EFI_TIME_CAPABILITIES   TimeCapabilities;

Status = gRT->GetTime (
                &Time,
                &TimeCapabilities
```

```
                 );

Status = gRT->GetTime (
                 &Time,
                 NULL       // Capabilities is optional and may be NULL
                 );
```

Example 4-35. Time and Date Services

4.5 EFI Driver Library

Table C-4 in Appendix C contains the list of EFI Driver Library Services that are available to EFI drivers. The details of these services can be found in the *EFI 1.10 Driver Library Specification*. These library services aid in the development of EFI drivers in several ways:

- These library services have been tested in a wide variety of driver types, so they are well validated.
- The EFI Driver Library is a lightweight library that helps reduce driver size compared to other EFI libraries.
- The EFI Driver Library is designed to complement the services that are already provided by the EFI Boot Services Table, the EFI Runtime Services Table, and the various protocol services. By using the combination of all of these services and the EFI Driver Library, the driver writer can concentrate on implementing the code that directly applies to the device being managed.

5
General Driver Design Guidelines

This chapter contains general guidelines for the design of all types of EFI drivers. Specific guidelines for PCI drivers, USB drivers, and SCSI drivers are presented in later chapters.

5.1 General Porting Considerations

Some EFI drivers are ported from IA-32 BIOS designs or older EFI 1.02 drivers. The process of porting a driver from one environment to another is often done to save time and leverage resources. If the port is done before obtaining a complete understanding of the target environment, the final driver may have remnants of the previous design that may not belong in the new environment. This section describes the practical porting issues (design and coding) that should be carefully reviewed before porting a driver.

5.1.1 How to Implement Features in EFI

The first column of Table 5-1 describes functions that a typical driver performs. The second column briefly describes how this function is implemented in EFI and references the section in this guide that specifically addresses each issue. This list of driver operations is not exhaustive.

Table 5-1. Mapping Operations to EFI Drivers

Operation	Recommended EFI Method
Find devices that the driver supports	Do not search. An EFI driver is passed a controller handle to evaluate along with a partial device path. See chapter 6.
Perform DMA	Use the DMA-related services from the PCI I/O Protocol. See section 14.5.
Access PCI configuration header	Use PCI I/O Protocol services. Never directly access I/O ports 0xCF8 or 0xCFC. See chapter 14 and section 17.2.
Access PCI I/O ports	Use PCI I/O Protocol services. Never use **IN** or **OUT** instructions. See chapter 14 and section 17.2.
Access PCI memory	Use PCI I/O Protocol services. Never use pointers to directly access memory-mapped I/O resources on a bus. See chapter 14 and section 17.2.
Hardware interrupts	EFI does not support hooking interrupts. Instead, EFI drivers are expected to either perform block I/O where they must complete their I/O operation and poll their device as required to complete it, or they can create a periodic timer event to allow the EFI driver to periodically get control to check the status of the devices it is managing. See section 4.4.2 and section 5.6.

continued

Table 5-1. Mapping Operations to EFI Drivers (continued)

Operation	Recommended EFI Method
Calibrated stalls	Do not use hardware devices to perform calibrated stalls. Instead, use the `gBS->Stall()` service for short delays that are typically less than 10 mS, and one-shot timer events for long delays that are typically greater than 10 mS. The `gRT->GetTime()` service should not be used for delays in EFI drivers. See section 4.4.1.
Get keyboard input from user	Use console services from the `SetOptions()` service of the Driver Configuration Protocol. Do not use console services anywhere else in the driver because console services may not be available when the driver executes. See chapter 11.
Display text	Use console services from the `SetOptions()` service of the Driver Configuration Protocol. Do not print messages from anywhere else in the driver because the console services may not be available when the driver executed. The `DEBUG()` macro can be used to send debug messages to the standard error console device. See chapter 11 and section 21.7.
Diagnostics	Implement the Driver Diagnostics Protocol. See chapter 12.
Flash utility	Write an EFI utility to reprogram the flash device.
Prepare controllers for use by an OS	The OS-present drivers should not make assumptions about the state of a controller. It should not assume that the controller was touched by an EFI driver before the OS was booted. If a specific state is required, then the driver can use an Exit Boot Services event to put the controller into the required state. See section 7.1.3.

5.2 Design and Implementation of EFI Drivers

The following is the list of basic steps that a driver writer should follow when designing and implementing an EFI driver.

1. Determine the class of EFI driver that needs to be developed. The different classes are listed in Table 5-3 below and are described in more detail in chapter 6 of this document.

Table 5-2. Classes of EFI Drivers to Develop

Class of Driver	For more information, see sections…
Device driver	6.1 and 9
Bus driver that can produce one or all child handles [Note 1]	6.2, 6.2.6, and 9
Bus driver that produces all child handles in the first call to `Start()`	6.2, 6.2.7, and 9
Bus driver that produces at most one child handle in `Start()`	6.2, 6.2.8, and 9

continued

Table 5-2. Classes of EFI Drivers to Develop (continued)

Class of Driver	For more information, see sections…
Bus driver that produces no child handles in `Start()`	6.2, 6.2.9, and 9
Bus driver that produces child handles with multiple parent controllers	6.2, 6.2.10, and 9
Hybrid driver that can produce one or all child handles [Note 2]	6.3, 6.2.6, and 9
Hybrid driver that produces all child handles in the first call to `Start()`	6.3, 6.2.7, and 9
Hybrid driver that produces at most one child handle in `Start()`	6.3, 6.2.8, and 9
Hybrid driver that produces no child handles in `Start()`	6.3, 6.2.9, and 9
Hybrid driver that produces child handles with multiple parent controllers	6.3, 6.2.10, and 9
Service driver	6.4 and 7.3
Root bridge driver	6.5 and 7.4
Initializing driver	6.6 and 7.2

Notes:

1 This bus driver type is recommended because it will enable faster boot times.

2 This hybrid driver type is recommended because it will enable faster boot times.

2. Is the EFI driver unloadable? The answer to this question is typically yes. See section 7.1.2 for details on how to enable this feature.
3. Is the EFI driver going to produce the Component Name Protocol? See chapter 10. This step is strongly recommended for all drivers.
4. Is the EFI driver going to produce the Driver Configuration Protocol? See chapter 11. This step is recommended.
5. Is the EFI driver going to produce the Driver Diagnostics Protocol? See chapter 12. This step is recommended.
6. If the EFI driver is a bus driver for a bus type that supports storage of EFI drivers with the child devices, then the Bus Specific Driver Override Protocol must be implemented by the bus driver. See chapter 13.
7. Is the EFI driver going to require an Exit Boot Services event? See section 7.1.3.
8. Is the EFI driver a runtime driver? See section 7.5 and section 20.1.
9. Is the EFI driver going to require a Set Virtual Address Map event? See section 7.5.
10. If the EFI driver follows the EFI Driver Model, then identify the I/O-related protocols that the driver needs to consume. Based on the list of consumed protocols and the criteria for these protocol interfaces, determine how many instances of the Driver Binding Protocol need to be produced. See sections 6.1.5 and 6.2.5.
11. If the EFI driver follows the EFI Driver Model, then identify the I/O-related protocols that the driver needs to produce. All device drivers, bus drivers, and hybrid drivers will use this method.
12. Implement the driver's entry point. See chapter 7.
13. Design the private context data structure. See chapter 8.
14. If the EFI driver follows the EFI Driver Model, then implement the `Supported()`, `Start()`, and `Stop()` services of each Driver Binding Protocol. See chapter 9.

15. Implement functions for each produced protocol. See the *EFI 1.10 Specification* for details on the protocols that are being produced.
16. Is the EFI driver for a PCI-related device? See chapter 14.
17. Is the EFI driver for a USB-related device? See chapter 15.
18. Is the EFI driver for a SCSI-related device? See chapter 16.
19. Is the EFI driver going to see a native driver for the Itanium processor? See chapter 18.
20. Is the EFI driver going to be an EBC driver? See chapter 19.
21. Design for maximum portability. See the porting considerations for Itanium architecture and EBC in chapter 18 and chapter 19.
22. Design for small code size and improved performance. See chapter 17.

5.3 Maximize Platform Compatibility

EFI drivers should make as few assumptions about a system's architecture as possible. Minimizing the number of assumptions will maximize the EFI driver's platform compatibility. It will also reduce the amount of driver maintenance when a driver is deployed on new systems.

5.3.1 Do Not Make Assumptions about System Memory Configurations

Do not make any assumptions about the system memory configuration, including memory allocations and memory that is used for DMA buffers. There may be unexpected gaps in the memory map, and entire memory regions may be missing.

EFI is designed for a wide variety of platforms. As such, portable drivers should not have artificial hard-coded limits; instead, they should rely on published specifications, EFI, and the system firmware to provide them with the platform limitations and platform resources, including the following:

- The number of adapters that can be supported in a system
- The type of adapter that can be supported on each bus
- The available memory resources

In addition, drivers should not make assumptions on a platform. Instead, they should make sure they support all the cases that are allowed by the *EFI 1.10 Specification*. For example, memory will not always be available beneath the 4 GB boundary (some systems may not have any memory under 4 GB at all) and drivers have to be designed to be compatible with these types of system configurations. As another example, some systems do not support PC-AT* legacy hardware and drivers should not expect them to be present.

The **gBS->AllocatePool()** service does not allow the caller to specify a preferred address, so this service is always safe to use and will have no impact on platform compatibility. The **gBS->AllocatePages()** service does have a mode that allows a specific address to be specified or a range of addresses to be specified. The allocation type of **AllocateAnyPages** is safe to use and will increase platform compatibility. The allocation types of **AllocateMaxAddress** and **AllocateAddress** may reduce platform compatibility.

An EFI driver should never directly allocate a memory buffer for DMA access. The EFI driver cannot know enough about the system architecture to predict what system memory areas are available for DMA. Instead, an EFI driver should use the services that are provided for the I/O bus to allocate and free buffers available for DMA. There should also be services to initiate and complete DMA transactions. For example, the PCI Root Bridge I/O Protocol and PCI I/O Protocol both provide services for PCI DMA operations. As additional I/O bus types with DMA capabilities are introduced, new protocols that abstract the DMA services will have to be added.

5.3.2 Do Not Use Any Hard-Coded Limits

EFI drivers should not use fixed-size arrays. Instead, memory resources should be dynamically allocated using the **gBS->AllocatePages()** and **gBS->AllocatePool()** services.

5.3.3 Do Not Make Assumptions about I/O Subsystem Configurations

EFI drivers should not assume there is a fixed or maximum number of controllers in a system. All EFI drivers that follow the EFI Driver Model should be designed to manage any number of controllers even if the driver writer is convinced there will only be a fixed number of controllers. This design will maximize the compatibility of the EFI driver, especially on multi-bus-set (ECR pending at PCI SIG) PCI systems that may contain hundreds of PCI slots. Chapter 8 introduces the private context data structure, which is a lightweight mechanism that allows an EFI driver to be designed with no limitations on the number of controllers that the EFI driver can manage.

5.3.4 Maximize Source Code Portability

EFI drivers should be designed to maximize source code portability. Today, the processor targets include the following:

- IA-32
- Itanium processor
- EBC virtual machine

It is possible to write a single driver that can be compiled for all three of these processor targets. It is also possible that the list of processor targets may grow over time. As a result, assembly language is discouraged.

An EFI driver should also never directly access any system chipset resources. Directly accessing these resources will limit the compatibility of the EFI driver to systems only with that specific chipset. Instead, the EFI Boot Services, EFI Runtime Services, and various protocol services should be used to access the system resources that are required by an EFI driver. Putting effort into source code portability will help maximize future platform compatibility.

5.4 EFI Driver Model

The EFI Driver Model goes a long way in addressing many of the potential platform differences and opens up a powerful way for drivers to interact with the user regardless of the platform specifics. The EFI Driver Model works with existing and future bus types. EFI drivers will typically be written to follow the EFI Driver Model.

The basic structure of an EFI driver that follows the EFI Driver Model is defined by several basic driver protocols. Each of the protocols adds standard interface functions that are coded in the EFI driver. An EFI driver needs to implement the following EFI Driver Model–related protocols and services:

- **`EFI_LOADED_IMAGE_PROTOCOL.Unload()`**
- **`EFI_DRIVER_BINDING_PROTOCOL.Supported()`**
- **`EFI_DRIVER_BINDING_PROTOCOL.Start()`**
- **`EFI_DRIVER_BINDING_PROTOCOL.Stop()`**
- **`EFI_DRIVER_CONFIGURATION_PROTOCOL.SetOptions()`**
- **`EFI_DRIVER_CONFIGURATION_PROTOCOL.OptionValid()`**
- **`EFI_DRIVER_CONFIGURATION_PROTOCOL.ForceDefaults()`**
- **`EFI_DRIVER_DIAGNOSTICS_PROTOCOL.RunDiagnostics()`**
- **`EFI_COMPONENT_NAME_PROTOCOL.GetDriverName()`**
- **`EFI_COMPONENT_NAME_PROTOCOL.GetControllerName()`**

5.5 Use the EFI Software Abstractions

EFI drivers shall use software abstractions that are provided by EFI and avoid the temptation of internal routines, unless EFI does not provide a suitable alternative. These software abstractions include the following:

- EFI Boot Services such as **`gBS->SetMem()`** and **`gBS->CopyMem()`**
- EFI Runtime Services such as **`gRT->GetVariable()`**
- EFI protocols such as the PCI I/O and Simple Text Output Protocols

This recommendation serves several purposes. By using the software abstractions provided by the platform vendor, the EFI driver will maximize its platform compatibility. The platform vendor can also optimize the services that are provided by the platform, so the performance of the EFI driver improves by using these services. Chapter 19 discusses the EBC porting considerations, and one of the most important considerations is the performance of an EBC driver because EBC code is interpreted. The performance of an EBC driver can be greatly improved by calling system services instead of using internal functions.

5.6 Use Polling Device Drivers

EFI drivers are not designed to be high-performance drivers, but rather to provide basic boot support for OS loaders. For this reason, EFI does not support an interrupt model for the device drivers. Instead, all EFI drivers operate in a polled mode. EFI drivers that implement blocking I/O services can simply poll the I/O device until the I/O request is complete. EFI drivers that implement

nonblocking I/O can create a periodic timer event to poll a device at periodic intervals. The interval should be set to the largest possible period for the EFI driver to complete its I/O services in a reasonable period of time. The overall performance of an EFI-enabled platform will degrade if too many EFI drivers create high-frequency periodic timer events. It is recommended that the period of a periodic timer event be at least 10 mS. In general, they should be as large as possible based upon a specific device's timing requirements, and most drivers can use events with timer periods in the range of 100 mS to several seconds. EFI drivers should also not spend a lot of time in their event notification functions, because this blocks the normal execution mode of the system. An EFI driver using a periodic timer event can always save some state information and wait for the next timer tick if the driver needs to wait for a device to respond. The ISA floppy driver and USB bus driver are examples of drivers in the *EFI 1.10 Sample Implementation* that use timer events.

5.6.1 Use Events and Task Priority Levels

The TPLs provide a mechanism for code to run at a higher priority than application code. So, one can be running the EFI Shell, and an EFI device driver can have a timer event fire and gain control to go poll its device. The **`TPL_CALLBACK`** level is typically used for deferred software calls, and **`TPL_NOTIFY`** is typically used by device drivers. **`TPL_HIGH_LEVEL`** is typically used for locks on shared data structures.

Drivers may use events and TPLs if they perform nonblocking I/O. If they perform blocking I/O, then they will not use events. They may still use the **`gBS->RaiseTPL()`** and **`gBS->RestoreTPL()`** for critical sections.

Driver diagnostics are typically just applications. They will not normally need to use TPLs or events unless the diagnostics is testing the TPL or event mechanisms in EFI. There is one exception. If a diagnostic needs to guarantee that EFI's timer interrupt is disabled, then the diagnostic can raise the TPL to **`TPL_HIGH_LEVEL`**. If this level is required, then it should be done for the shortest possible time interval.

5.7 Design to Be Re-entrant

If a system contains multiple controllers of the same type from the same vendor, it is quite possible that a single driver could install an instance of the controller's I/O protocol on a handle for each device. Each instance of the protocol would call the same driver functions, but the data would be unique to the instance (or context) of the protocol. This design concept is important for all EFI drivers.

The practical manifestation of this requirement is that all the data that must be local to the instance (context) of the protocol must **not** be stored in global variables. Instead, data is collected into a private context data structure and each time that an I/O protocol is installed onto a handle, a new version of the structure is allocated from memory. This concept is described in detail in chapter 8.

5.8 Avoid Function Name Collisions between Drivers

Compilers and linkers will guarantee that there are no name collisions within a driver, but the compilers and linkers cannot check for name collisions between drivers. This inability to check is a concern only when debuggers are used that can perform source-level debugging or can display function names. Section 3.12 introduced templates that help avoid function name collisions between drivers. The guidelines in the templates should be followed.

5.9 Manage Memory Ordering Issues in the DMA and Processor

Not all processors have strongly ordered memory models. This distinction means that the order in which memory transactions are presented in the source code may not be the same when the code is executed. Normally, this order is not an issue, because the processor and the compiler will guarantee that the code will execute as the developer expects. However, EFI drivers that use DMA buffers that are simultaneously accessed by both the processor and the DMA bus master may run into issues if either the processor or the DMA bus master or both are weakly ordered. The DMA bus master must solve its own ordering issues, but the EFI driver writer is responsible for managing the processor's ordering issues.

The *EFI 1.10 Sample Implementation* contains a macro called **`MEMORY_FENCE()`** that guarantees that all the transactions in the source code prior to the **`MEMORY_FENCE()`** macro are completed before the code after the **`MEMORY_FENCE()`** macro is executed. This macro can be used in an EFI driver that accesses a buffer at the same time that a DMA bus master is accessing that same buffer. The EFI driver will insert these macros at the appropriate points in the code where all the processor transactions need to be completed before the next processor transaction is performed. For example, if a buffer contains a data structure and that data structure contains a valid bit and additional data fields that describe an I/O operation that the DMA bus master needs to perform, then the EFI driver would want to update the additional data fields before setting the valid bit. The EFI driver would use the **`MEMORY_FENCE()`** macro just before and just after the valid bit is written to the buffer.

This issue is not present in IA-32 or EBC virtual machines because these two processor types are strongly ordered. However, Itanium-based platforms are weakly ordered, so this macro must be used for native drivers for the Itanium processor. It is recommended that these macros be used appropriately in all driver types to maximize the EFI driver's platform compatibility.

5.10 Do Not Store EFI Drivers in Hidden Option ROM Regions

Some option ROMs may use paging or other techniques to load and execute code that was not visible to the system firmware when measuring the visible portion of the option ROM. This technique is discouraged because it is typically the bus driver's responsibility to extract the option ROM contents when a bus is enumerated. If code is required to access hidden portions of an option ROM, then the bus driver would not have the ability to extract the additional option ROM contents. This inability means that the EFI drivers in an option ROM must be visible without accessing a hidden portion of an option ROM. However, if there is a safe mechanism to access the hidden portions of the option ROM after the EFI drivers have been loaded and executed, then the EFI

driver may choose to access those contents. For example, nonvolatile configuration information, utilities, or diagnostics can be stored in the hidden option ROM regions.

5.11 Store Configuration Settings on the Same FRU as the EFI Driver

The configuration for an EFI driver should be stored on the same Field Replaceable Unit (FRU) as the EFI driver. If an EFI driver is stored on the motherboard, then their configuration information can be stored in EFI variables. If an EFI driver is stored in an add-in card, then their configuration information should be stored in the NVRAM provided on the add-in card.

5.11.1 Benefits

This method ensures that it is possible to statically (while the FRU is being designed) determine the maximum configuration storage that is required for the FRU. In particular, if option cards were to store their configuration in EFI variables, the amount of variable storage could not be statically calculated, because it generally is not possible to know ahead of time what set of option cards may be installed in a system. The result would be that add-in cards could not be used in otherwise functional systems due to lack of EFI variable storage space.

Storing EFI variables in the same FRU as the EFI driver reduces the amount of stale data that is left in EFI variables. If an option card was to store its data in EFI variables and then be removed, there is no automatic cleanup mechanism to purge the EFI variables that are associated with that card.

Storing EFI variables in the same FRU as the EFI driver also ensures that the configuration stays with the FRU. It enables centralized configuration of add-in cards. For example, if an IT department is configuring 50 like systems, it can configure all 50 in the same system and then disburse them to the systems, rather than configuring each system separately. Further, it can maintain preconfigured spares.

5.11.2 Update Configurations at OS Runtime Using an OS-Present Driver

If an end-user or developer wants to configure an add-in card while the OS is running, it is possible to update the configuration to the card using EFI variables. Add-in cards are typically supplied with OS-present drivers. For most operating systems, it is actually preferable to access the card using the OS-driver and for that driver to update the card.

5.12 Do Not Use Hard-Coded Device Path Nodes

The **`ACPI()`** node in the EFI Device Path Protocol identifies the PCI root bridge in the ACPI namespace. The *ACPI Specification* allows _HID to describe vendor-specific capability and _CID to describe compatibility. Therefore, there is no requirement for all platforms to use the PNP0A03 identifier in the _HID to identify the PCI root bridge. The following are the only requirements for the PCI root bridge:

- The PNP0A03 identifier must appear in _HID if a vendor-specific capability does not need to be described.
- The PNP0A03 identifier must appear in _CID if _HID contains a vendor-specific identifier.

To avoid problems with platform differences, EFI drivers should not create EFI device paths from hard-coded information. Instead, EFI bus drivers should append new device path nodes to the device path from the parent device handle.

5.12.1 PNPID Byte Order for EFI

The ACPI PNPID format (byte order) follows the original EISA ID format. EFI also uses PNPID in the device path ACPI nodes. However, for a given string, ACPI and EFI do not generate the same numbers. For example:

```
HID = "PNP0501"
ACPI = 0x0105D041
EFI = 0x050141D0
```

This difference means that operating systems that try to match the EFI ACPI device path node to the ACPI name space must perform a translation.

5.13 Do Not Cause Errors on Shared Storage Devices

In a cluster configuration, multiple devices may be connected to a shared storage. In such configurations, the EFI driver should not cause errors that can be seen by the other devices that are connected to storage.

On a boot or reboot, there shall be no writes to shared storage without user acknowledgement. Any writes to shared storage by an EFI driver may corrupt shared storage as viewed by another system. As a result, all outstanding I/O in the controller's buffers shall be cleared, and any internal caches shall also be cleared. Any I/O operations that occur after a reboot may corrupt shared storage.

There must not be an excessive number of bus or device resets. Device resets have an impact on shared storage as viewed by other systems. For a single reset, this impact is negligible. Larger numbers of resets may be seen as a device failure by another system.

Disk signatures must not be changed without warning the user. If there is an impact to the user, then that impact should be displayed along with the warning. Clusters may make an assumption about disk signatures on shared storage.

The discovery process must not impact other systems accessing the storage. A long discovery process may "hold" drives and look like a failure of shared storage.

5.14 Convert Bus Walks

In EFI, the EFI bus drivers find, enumerate, and expose devices using the bus's standard I/O protocol. For example, the PCI bus driver exposes devices on the PCI bus that have device handles that support the PCI I/O Protocol. The USB bus driver exposes USB devices that have a handle that supports the USB I/O Protocol. The list of device handles that support a specific bus I/O protocol can be discovered using the **gBS->LocateHandleBuffer()** service. Not only is this method faster than a classical bus walk, but it will work even when a system contains multiple PCI buses.

5.14.1 Example

Example 5-1 shows an example of converting a bus walk on the PCI bus. In this example, use EFI Boot Services to obtain all the handles that support the PCI I/O Protocol and then examine the configuration space for the vendor ID (VID) and device ID (DID). Note that the code below could be used in a flash utility to discover all the cards in the system.

```
EFI_STATUS              Status;
UINTN                   HandleCount;
EFI_HANDLE              *HandleBuffer;
UINTN                   Index;
EFI_PCI_IO_PROTOCOL     *PciIo;
PCI_TYPE00              Pci;

//
// Retrieve the list of handles that support the PCI I/O Protocol from
// the handle database. The number of handles that support the PCI I/O
// Protocol is returned in HandleCount, and the array of handle values is
// returned in HandleBuffer.
//
Status = gBS->LocateHandleBuffer (
                ByProtocol,
                &gEfiPciIoProtocolGuid,
                NULL,
                &HandleCount,
                &HandleBuffer
                );
if (EFI_ERROR (Status)) {
  return Status;
}

//
// Loop through all the handles that support the PCI I/O Protocol, and
// retrieve the instance of the PCI I/O Protocol.
//
for (Index = 0; Index < HandleCount; Index++) {
  Status = gBS->OpenProtocol (
                  HandleBuffer[Index],
                  &gEfiPciIoProtocolGuid,
                  (VOID **)&PciIo,
                  ImageHandle,
                  NULL,
                  EFI_OPEN_PROTOCOL_GET_PROTOCOL
                  );
  if (EFI_ERROR (Status)) {
    continue;
  }
```

```
  DEBUG ((D_INIT, "DrvUtilApp: Call Pci.Read \n"));
  Status = PciIo->Pci.Read (
                        PciIo,
                        EfiPciIoWidthUint8,
                        0,
                        sizeof (Pci),
                        &Pci
                        );

  if (EFI_ERROR (Status)) {
    Status = EFI_UNSUPPORTED;
    DEBUG ((D_INIT, "DrvUtilApp: Return Value = %r\n",Status));
  } else {

    //
    // Check for desired attributes:
    //  Pci.Hdr.VendorId
    //  Pci.Hdr.DeviceId
    //

    if ((Pci.Hdr.VendorId != XYZ_VENDOR_ID) ||
        (Pci.Hdr.DeviceId != XYZ_DEVICE_ID)) {
      DEBUG((D_INIT, "DrvUtilApp: This DH not ours: %d\n", HandleBuffer[Index]));
    } else {
      DEBUG((D_INIT, "DrvUtilApp: FOUND! %d\n", HandleBuffer[Index]));
    }
  }
}
//
// Free the array of handles that was allocated by gBS->LocateHandleBuffer()
//
gBS->FreePool (HandleBuffer);
```

Example 5-1. PCI Bus Walk Example

5.15 Do Not Have Any Console Display or Hot Keys

PC BIOS legacy option ROMs would typically display banners. The header displays the driver name and version information. Some BIOS drivers inform the user of a "hot key" that would be checked during the boot sequence that would allow customers to enter into the driver's configuration options. This type of interaction created several problems:

- Displaying the banner and waiting long enough to detect a key press increases boot time. On systems where there could be hundreds of PCI cards, this time can add up. Even the time to simply print a header can take too long, because some systems may require the console text to flow out to a serial port that could be configured at 9600 baud.
- The user is rarely ready to interact with the boot driver at just the right moment and often has to reboot the system multiple times to read the display banner or press the key at the right time to enter the driver configuration utility. On larger systems, the system boot can take many minutes.
- The amount of text "spew" can be confusing to a user watching the boot sequence. Each driver can have its own unique way of presenting the information. When many cards are all displaying their own unique header information to the screen at once, there is a significant risk of

confusing the user with too much information that is only informational and not significant to the health of the system. Also, the text can unnecessarily fill console logs and potentially break scripts when a new card is added to a system that was monitoring the console text.

- On fast systems, the header text is displayed but erased by a later driver so quickly that the user cannot read it.
- In EFI, the driver may be initialized before the console output and standard error devices. In such cases, the data will not be visible or will crash the system. When EFI option ROM drivers are loaded, the console has not been connected because the console may in fact be controlled by one of the EFI option ROM drivers. In such cases, printing may in fact call a **NULL** function and crash the system.

For these reasons, the EFI driver model requires that no console I/O operations take place in the EFI Driver Binding Protocol functions. A reasonable exception to this rule is to use the **DEBUG()** macro to display progress information during driver development and debug. Using the **DEBUG()** macro allows the code for displaying the data to be easily removed for a production build of the driver.

Use of the **DEBUG()** macro should be limited to "debug releases" of a driver. This strategy will typically work if the driver is loaded after the EFI console is connected. However, some firmware implementations may load the option ROM drivers before the EFI console is connected (because console drivers may live in option ROMs). In such cases, the *ConOut* and *StdErr* fields of the EFI system table may be **NULL**, and printing will crash the system. The **DEBUG()** macro should check to see if the field is **NULL** before using those services.

> **Tip:** The *EFI 1.10.14.59 Sample Implementation* fixed a bug in the print function that is used by EFI drivers so that print format strings using **%s** would print a Unicode string rather than an ASCII string as it did earlier. To avoid problems on systems that may not have the fix, consider using **%a** or **%S** instead of **%s**, where **%S** prints a Unicode string and **%a** prints an ASCII string.

Chapter 9 of the *EFI 1.10 Specification* provides several other protocols to assist in interacting with the user—the EFI Driver Configuration Protocol and EFI Driver Diagnostics Protocol.

5.16 Offer Alternatives to Function Key

The EFI console may be connected through a serial port. In such cases, it is sensitive to the correct terminal emulator configuration. If the user has not correctly configured the terminal emulator to match the terminal settings in EFI (PC ANSI, VT100, VT100+, or VT-UTF8), they may not be able to correctly use some of the keys (function keys, arrow keys, page up/down, insert/delete, and backspace) or be able to display colors and see the correct cursor positioning. To better support users, it is recommended that EFI configuration protocols and EFI applications create user interfaces that are not solely dependent on these keys but instead offer alternatives for these keys. Also, it is important to note that the Simple Input Protocol does not support the CTRL or ALT keys because these keys are not available with remote terminals such as terminal emulators and telnet. Table 5-3 below shows one possible set of alternate key sequences for function keys, arrow keys, page up/down keys, and the insert/delete keys. Each configuration protocol and application will have to decide if alternate key sequences are supported and which alternate mappings should be used. Table 5-3 lists the EFI Scan Code from the Simple Input Protocol and the alternate key sequence that can be used to produce that scan code. Most of these key sequences are directly

supported in the *EFI 1.10.14.62 Sample Implementation,* which means that the developer does not need to do anything special to support these key sequences on a remote terminal. The ones labeled as "No" are not directly supported in the EFI *1.10.14.62 Sample Implementation*, so those key sequences will have to be parsed to be interpreted by the configuration protocol or application.

Table 5-3. Alternate Key Sequences for Remote Terminals

EFI Scan Code	Key Sequence	Supported in EFI 1.10.14.62?
SCAN_NULL		
SCAN_UP	'^'	No
SCAN_DOWN	'v' or 'V'	No
SCAN_RIGHT	'>'	No
SCAN_LEFT	'<'	No
SCAN_HOME	ESC h	Yes
SCAN_END	ESC k	Yes
SCAN_INSERT	ESC +	Yes
SCAN_DELETE	ESC -	Yes
SCAN_PAGE_UP	ESC ?	Yes
SCAN_PAGE_DOWN	ESC /	Yes
SCAN_F1	ESC 1	Yes
SCAN_F2	ESC 2	Yes
SCAN_F3	ESC 3	Yes
SCAN_F4	ESC 4	Yes
SCAN_F5	ESC 5	Yes
SCAN_F6	ESC 6	Yes
SCAN_F7	ESC 7	Yes
SCAN_F8	ESC 8	Yes
SCAN_F9	ESC 9	Yes
SCAN_F10	ESC 0	Yes
ESC	ESC	Yes

5.17 Do Not Assume EFI Will Execute All Drivers

Typically, the same vendor that produces an EFI driver will also have to produce an OS-present driver for all the operating systems that the vendor chooses to support. Because EFI provides a mechanism to reduce the boot time by running the minimum set of drivers that are required to connect the console and boot devices, not all EFI drivers may be executed on every boot. For example, the system may have three SCSI cards but only needs to install the driver on one SCSI bus to boot the OS.

This minimum set of drivers means that the OS-present driver may be handed a controller that may be in several different states. It may still be in the power-on reset state, it may have been managed by an EFI driver for a short period of time and released, and it may have been managed by an EFI driver right up to the point in time where firmware hands control of the platform to the operating system.

The OS-present driver must accept controllers in all of these states. This acceptance requires the OS-present driver to make very few assumptions about the state of the controller it manages.

OS drivers shall not make assumptions that the EFI driver has initialized or configured the device in any way. Also, I/O hot-plug does not involve EFI driver execution, so the OS driver must be able to initialize and operate the driver without EFI support.

6
Classes of EFI Drivers

The different classes of EFI drivers are introduced in chapter 2. These driver classes will be discussed throughout this document, but an emphasis is placed on drivers that follow the EFI Driver Model because these drivers are the most commonly implemented. The drivers that follow the EFI Driver Model include the following:

- Device drivers
- Bus drivers
- Hybrid drivers

There are actually several subtypes and optional features for these classes of drivers. This chapter introduces the subtypes and optional features of drivers that follow the EFI Driver Model. Understanding the different classes of EFI drivers will help driver writers identify the class of driver to implement and the algorithms that are used in their implementation. The less common service drivers, root bridge drivers, and initializing drivers are also discussed.

The chapters that follow describe a driver's entry point in detail, including the various optional features that may be enabled in the entry point. The driver entry point chapter is followed by a chapter on using object-oriented programming techniques to help design good data structures in drivers that follow the EFI Driver Model. This chapter is followed by a chapter on the Driver Binding Protocol and chapters on the optional protocols such as Component Name, Driver Configuration, Driver Diagnostics, Bus Specific Driver Override, and Platform Driver Override.

6.1 Device Drivers

All device drivers that follow the EFI Driver Model share a set of common characteristics. The following two sections describe the required and optional features for device drivers. These sections are followed by a detailed description of device drivers that produce a single instance of the Driver Binding Protocol and device drivers that produce multiple instances of the Driver Binding Protocol.

6.1.1 Required Device Driver Features

Device drivers are required to implement the following features:

- Installs one or more instances of the **EFI_DRIVER_BINDING_PROTOCOL** in the driver's entry point.
- Manages one or more controller handles. Even if a driver writer is convinced that the driver will manage only a single controller, the driver should be designed to manage multiple controllers. The overhead for this functionality is low, and it will make the driver more portable.
- Does not produce any child handles. This feature is the main distinction between device drivers and bus drivers.
- Ignores the *RemainingDevicePath* parameter that is passed into the **Supported()** and **Start()** services of **EFI_DRIVER_BINDING_PROTOCOL**.
- Consumes one or more I/O-related protocols from a controller handle.
- Produces one or more I/O-related protocols on the same controller handle.

6.1.2 Optional Device Driver Features

The following is the list of features that device drivers can optionally implement:

- Installs one or more instances of the **EFI_COMPONENT_NAME_PROTOCOL** in the driver's entry point. Implementing this feature is strongly recommended. It allows a driver to provide human-readable names for the name of the driver and the controllers that the driver is managing. Some platform design guides such as *DIG64* require this feature.
- Installs one or more instances of the **EFI_DRIVER_CONFIGURATION_PROTOCOL** in the driver's entry point. If a driver has any configurable options, then this protocol is required. Some platform design guides such as *DIG64* require this feature.
- Installs one or more instances of the **EFI_DRIVER_DIAGNOSTICS_PROTOCOL** in the driver's entry point. If a driver wishes to provide diagnostics for the controllers that the driver manages, then this protocol is required. This protocol is the only mechanism that is available to a driver when the driver wants to alert the user of a problem that was detected with a controller. Some platform design guides such as *DIG64* require this feature.
- Provides an **EFI_LOADED_IMAGE_PROTOCOL.Unload()** service, so the driver can be dynamically unloaded. It is recommended that this feature be implemented during driver development, driver debug, and system integration. It is strongly recommended that this service remain in drivers for add-in adapters to help debug interaction issues during system integration.
- Creates an Exit Boot Services event in the driver's entry point. This feature is required only if the driver is required to place the devices it manages in a specific state just before control is handed to an operating system.
- Creates a Set Virtual Address Map event in the driver's entry point. This feature is required only for a device driver that is also an EFI runtime driver.

6.1.3 Device Drivers with One Driver Binding Protocol

Most device drivers produce a single instance of the **EFI_DRIVER_BINDING_PROTOCOL**. These drivers are the simplest that follow the EFI Driver Model, and all other driver types have their roots in this type of device driver.

A device driver is loaded into memory with the **gBS->LoadImage()** Boot Service and invoked with the **gBS->StartImage()** Boot Service. The **gBS->LoadImage()** service automatically creates an image handle and installs the **EFI_LOADED_IMAGE_PROTOCOL** onto the image handle. The **EFI_LOADED_IMAGE_PROTOCOL** describes the location where the device driver was loaded and the location in system memory where the device driver was placed. The **Unload()** service of the **EFI_LOADED_IMAGE_PROTOCOL** is initialized to **NULL** by **gBS->LoadImage()**. This setting means that by default the driver is not unloadable. The **gBS->StartImage()** service transfers control to the driver's entry point as described in the PE/COFF header of the driver image. The driver entry point is responsible for installing the **EFI_DRIVER_BINDING_PROTOCOL** onto the driver's image handle. Figure 6-1 below shows the state of the system before a device driver is loaded, just before it is started, and after the driver's entry point has been executed.

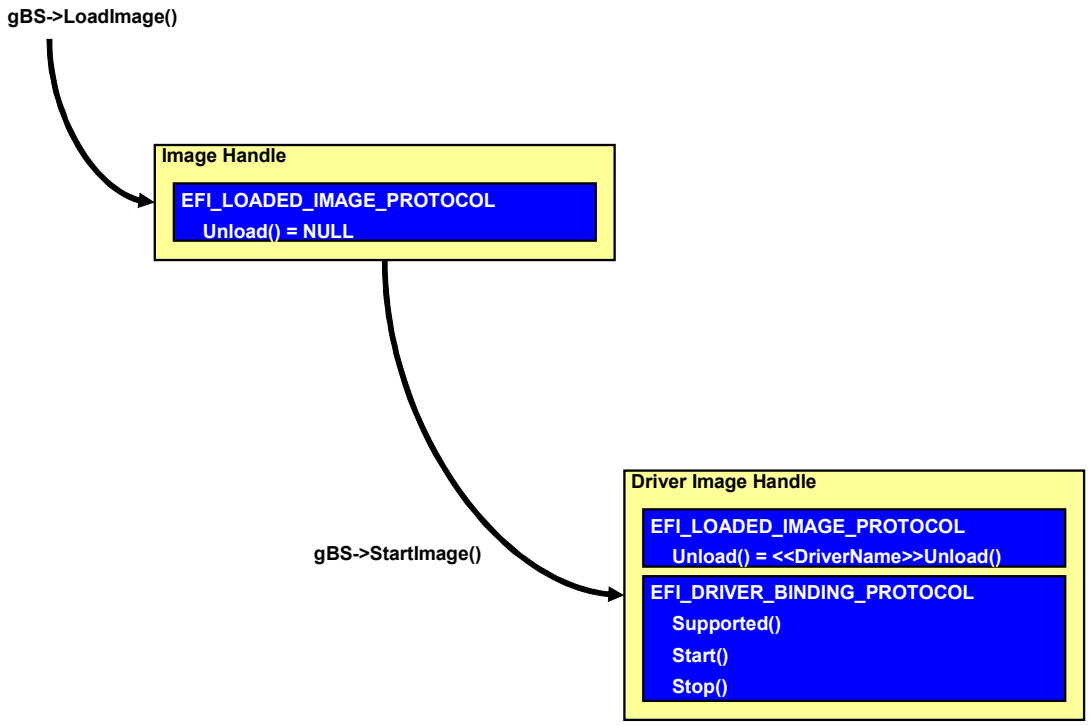

Figure 6-1. Device Driver with Single Driver Binding Protocol

Figure 6-2 below is the same as the figure above, except this device driver has also implemented all the optional features. This difference means the following:

- Additional protocols are installed onto the driver's image handle.
- An **Unload()** service is registered in the **EFI_LOADED_IMAGE_PROTOCOL**.
- An Exit Boot Services event and Set Virtual Address Map event have been created.

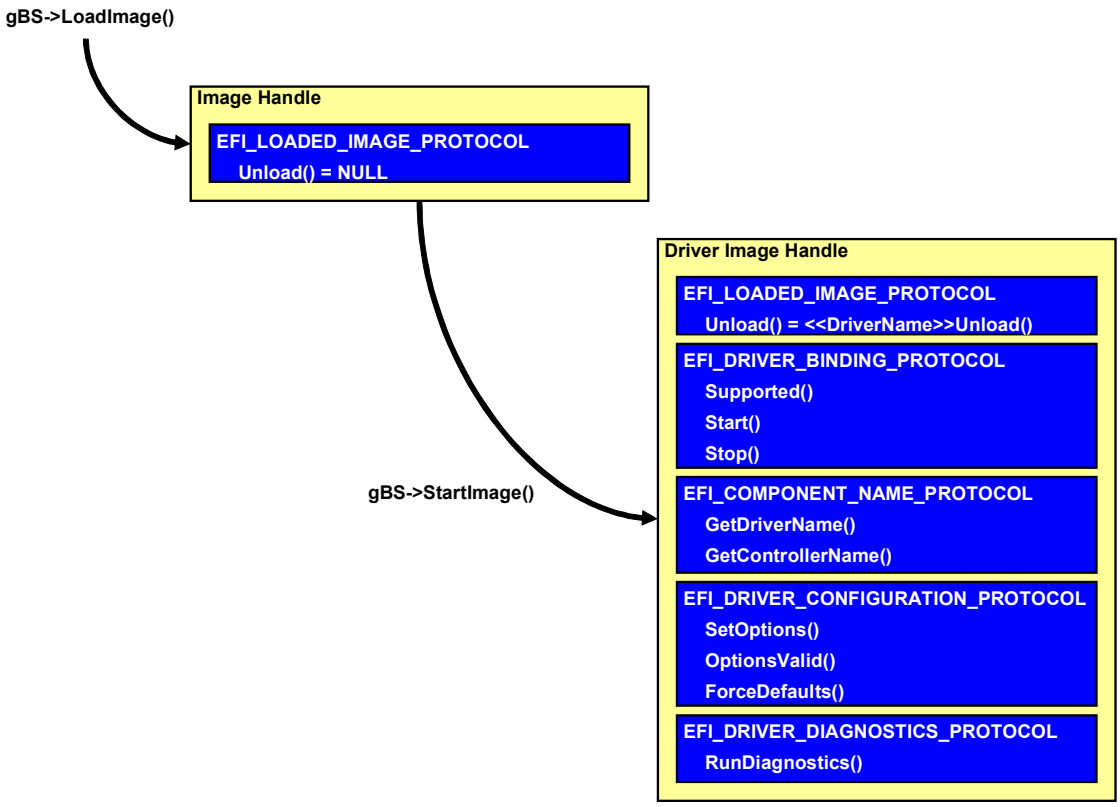

Figure 6-2. Device Driver with Optional Features

Table 6-1 below lists the device drivers from the *EFI 1.10 Sample Implementation* that produce a single instance of the **EFI_DRIVER_BINDING_PROTOCOL**.

Table 6-1. Device Drivers with One Driver Binding Protocol

AtapiPassThru	PciVgaMiniPort	Usb\UsbBot	WinNtThunk\SimpleFileSystem
CirrusLogic5430	Ps2Keyboard	Usb\UsbCbi	WinNtThunk\Uga
Console\GraphicsConsole	Ps2Mouse	Usb\UsbKb	BiosInt\BiosKeyboard
DiskIo	PxeBc	Usb\UsbMassStorage	BiosInt\BiosVga
FileSystem\Fat	PxeDhcp4	Usb\UsbMouse	BiosInt\BiosVgaMiniPort
IsaFloppy	ScsiDisk	VgaClass	
PcatIsaAcpi	Snp32_64	WinNtThunk\BlockIo	
PcatIsaAcpiBios	Usb\Uhci	WinNtThunk\Console	

6.1.4 Device Drivers with Multiple Driver Binding Protocols

A more complex device driver is one that produces more than one instance of the **EFI_DRIVER_BINDING_PROTOCOL**. The first instance of the **EFI_DRIVER_BINDING_PROTOCOL** is installed onto the driver's image handle, and the additional instances of the **EFI_DRIVER_BINDING_PROTOCOL** are installed onto newly created driver binding handles.

Figure 6-3 below shows the state of the handle database before a driver is loaded, before it is started, and after its driver entry point has been executed. This specific driver produces three instances of the **EFI_DRIVER_BINDING_PROTOCOL**.

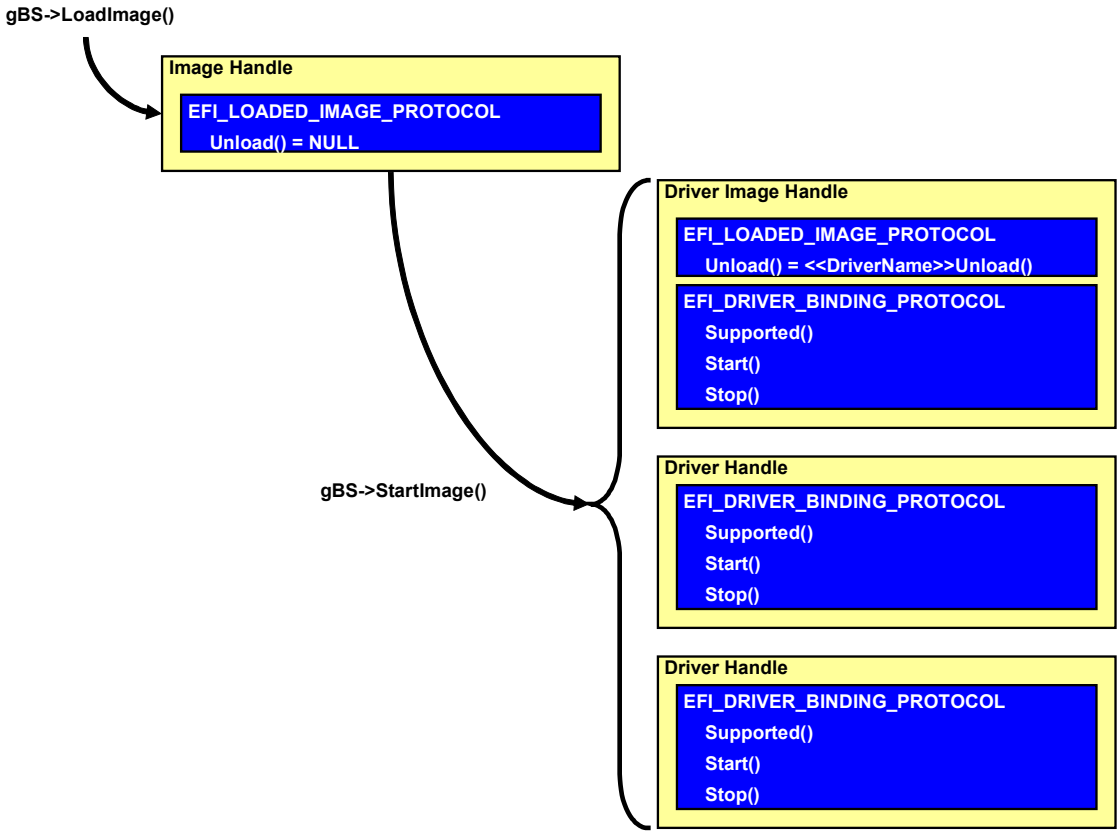

Figure 6-3. Device Driver with Multiple Driver Binding Protocols

Any device driver that produces multiple instances of the **EFI_DRIVER_BINDING_PROTOCOL** can be broken up into multiple drivers, so each driver produces a single instance of the **EFI_DRIVER_BINDING_PROTOCOL**. However, there are a few advantages for a driver to produce multiple instances of the **EFI_DRIVER_BINDING_PROTOCOL**. The first is that it may reduce the overall size of the drivers. If two related drivers are combined and those two drivers can share internal functions, the executable image size of the single driver may be smaller than the sum of the two individual drivers. The other reason to combine drivers is to help manage platform features. A single platform features may require several drivers. If the drivers are separated, then multiple drivers have to be added or removed to add or remove that single feature.

There is only one device driver in the *EFI 1.10 Sample Implementation* that produces multiple instances of the **EFI_DRIVER_BINDING_PROTOCOL**, and it is the console platform driver in the **\Efi1.1\Edk\Console\ConPlatform** directory. This driver implements the platform policy for managing multiple console input and console out devices. It produces one **EFI_DRIVER_BINDING_PROTOCOL** for the console output devices, and another **EFI_DRIVER_BINDING_PROTOCOL** for the console input devices. The management of console devices needs to be centralized, so it makes sense to combine these two functions into a single driver so the platform vendor needs to update only one driver to adjust the platform policy for managing console devices.

6.1.5 Device Driver Protocol Management

Device drivers consume one or more I/O-related protocols and use the services of those protocols to produce one or more I/O-related protocols. The **Supported()** and **Start()** functions of the **EFI_DRIVER_BINDING_PROTOCOL** are responsible for opening the I/O-related protocols that are being consumed using the EFI Boot Service **gBS->OpenProtocol()**. The **Stop()** function is responsible for closing the consumed I/O-related protocols using **gBS->CloseProtocol()**. A protocol can be opened in several different modes, but the most common is **BY_DRIVER**. When a protocol is opened **BY_DRIVER**, a test is made to see if that protocol is already being consumed by any other drivers. The open operation will succeed only if the protocol is not being consumed by any other drivers. This use of the **gBS->OpenProtocol()** service is how resource conflicts are avoided in the EFI Driver Model. However, it requires that every driver present in the system to follow the driver interoperability rules for all resource conflicts to be avoided.

Figure 6-4 below shows the image handle for a device driver as **gBS->LoadImage()** and **gBS->StartImage()** are called. In addition, it shows the states of three different controller handles as the **EFI_DRIVER_BINDING_PROTOCOL** services **Supported()**, **Start()**, and **Stop()** are called. **Controller Handle 1** and **Controller Handle 3** pass the **Supported()** test, so the **Start()** function can be called. In this case, the **Supported()** service tests to see if the controller handle supports Protocol A. **Start()** is then called for **Controller Handle 1** and **Controller Handle 3**. In the **Start()** function, **Protocol A** is opened **BY_DRIVER**, and **Protocol B** is installed onto the same controller handle. The implementation of **Protocol B** will use the services of **Protocol A** to produce the services of **Protocol B**. All drivers that follow the EFI Driver Model must support the **Stop()** service. The **Stop()** service must put the handles back into the same state they were in before **Start()** was called, so the **Stop()** service uninstalls **Protocol B** and closes **Protocol A**.

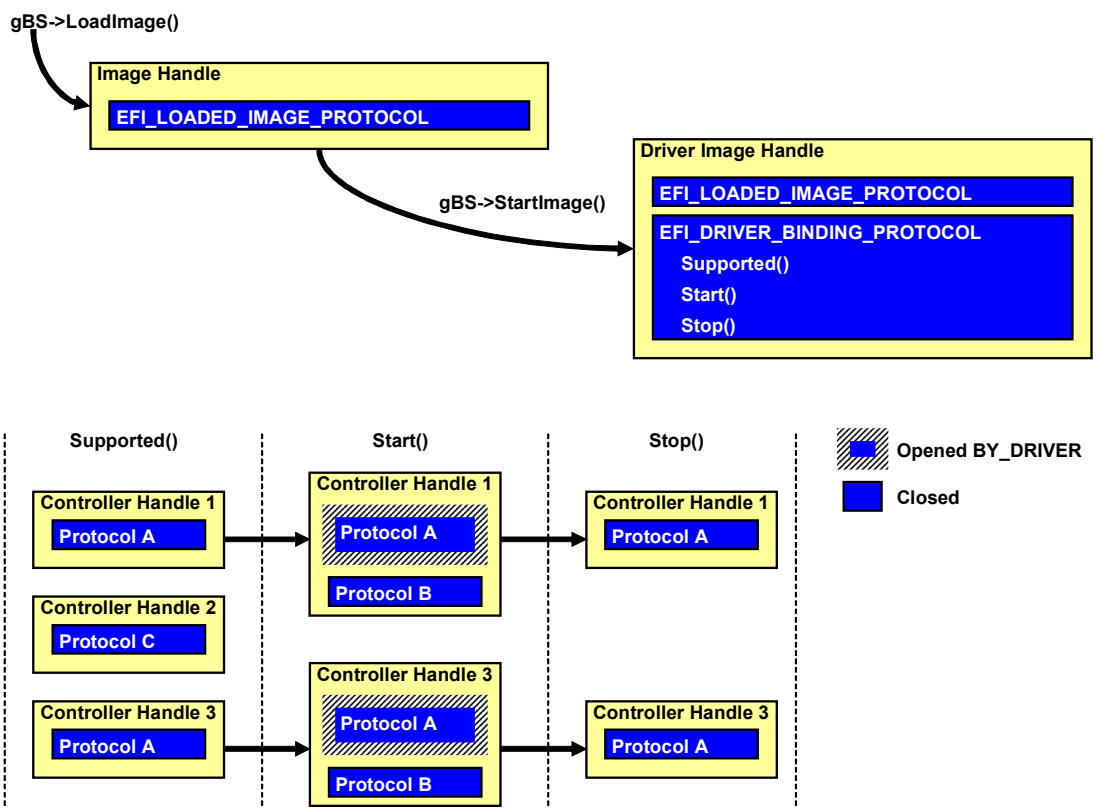

Figure 6-4. Device Driver Protocol Management

Figure 6-5 below shows a more complex device driver that requires **Protocol A** and **Protocol B** to produce **Protocol C**. Notice that the controller handles that support neither **Protocol A** nor **Protocol B**, only **Protocol A**, or only **Protocol B** do not pass the `Supported()` test. Also notice that **Controller Handle 6** already has **Protocol A** opened `BY_DRIVER`, so this device driver that requires both Protocol A and Protocol B will not pass the `Supported()` test either. This example highlights some of the flexibility of the EFI Driver Model. Because the `Supported()` and `Start()` services are functions, a driver writer can implement simple or complex algorithms to test if the driver supports a specific controller handle.

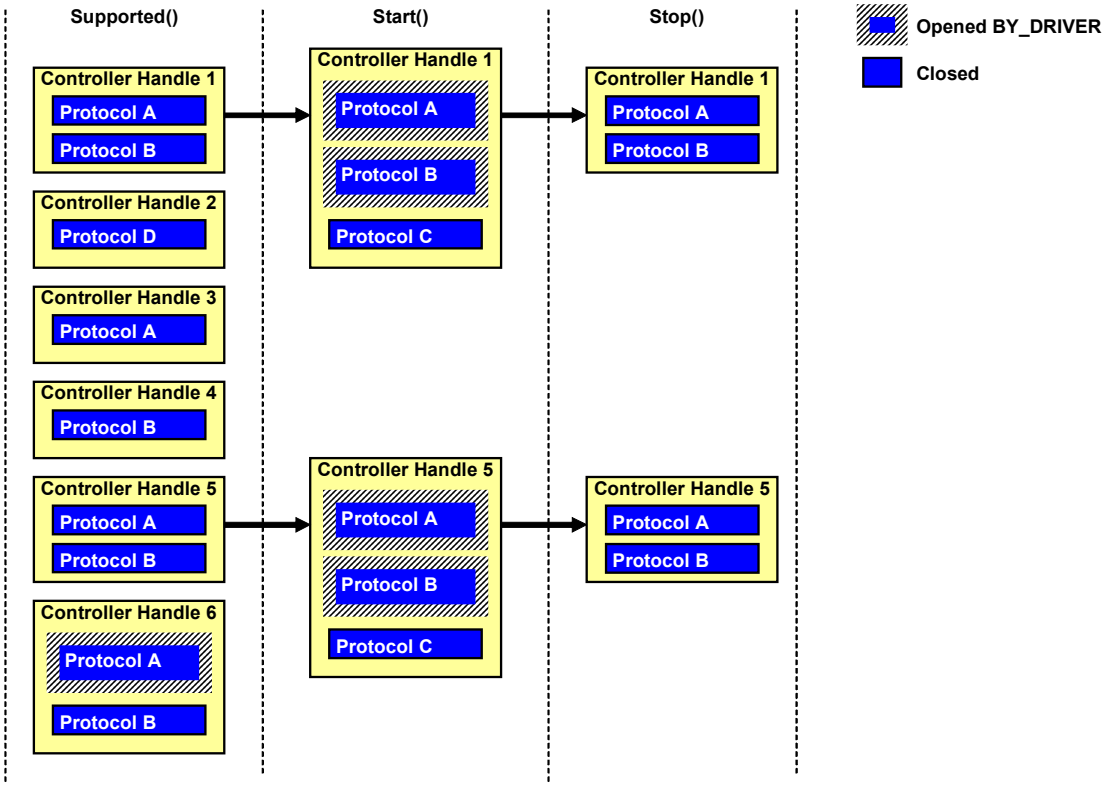

Figure 6-5. Complex Device Driver Protocol Management

The best way to design the algorithm for the opening protocols is to write a Boolean expression for the protocols that a device driver consumes. Then, expand this Boolean expression into the sum of products form. Each product in the expanded expression requires its own **`EFI_DRIVER_BINDING_PROTOCOL`**. This scenario is another way that a device driver may be required to produce multiple instances of the **`EFI_DRIVER_BINDING_PROTOCOL`**. The **`Supported()`** service for each **`EFI_DRIVER_BINDING_PROTOCOL`** will attempt to open each protocol in a product term. If any of those open operations fail, then **`Supported()`** fails. If all the opens succeed, then the **`Supported()`** test passes. The **`Start()`** function should open each protocol in the product term, and the **`Stop()`** function should close each protocol in the product term.

For example, the two examples above would have the following Boolean expressions:

(**Protocol A**)

(**Protocol A** AND **Protocol B**)

These two expressions have only one product term, so only one **EFI_DRIVER_BINDING_PROTOCOL** is required. A more complex expression would be as follows:

(**Protocol A** AND (**Protocol B** OR **Protocol C**))

If this Boolean expression is expanded into a sum of product form, it would yield the following:

((**Protocol A** AND **Protocol C**) OR (**Protocol B** AND **Protocol C**))

This expression would require a driver with two instances of the **EFI_DRIVER_BINDING_PROTOCOL**. One would test for **Protocol A** and **Protocol C**, and the other would test for **Protocol B** and **Protocol C**.

6.2 Bus Drivers

All bus drivers that follow the EFI Driver Model share a set of common characteristics. The following two sections describe the required and optional features for bus drivers. These sections are followed by a detailed description of bus drivers that do the following:

- Produce a single instance of the Driver Binding Protocol
- Produce multiple instances of the Driver Binding Protocol
- Produce all of their child devices in their **Start()** function
- Are able to produce a single child device in their **Start()** function
- Produce at most one child device from their **Start()** function
- Bus drivers for hot-plug bus types that do not produce any child devices in their **Start()** function
- Produce child devices with multiple parent devices

6.2.1 Required Bus Driver Features

Bus drivers are required to implement the following features:

- Installs one or more instances of the **EFI_DRIVER_BINDING_PROTOCOL** in the driver's entry point.
- Manages one or more controller handles. Even if a driver writer is convinced that the driver will manage only a single bus controller, the driver should be designed to manage multiple bus controllers. The overhead for this functionality is low, and it will make the driver more portable.
- Produces any child handles. This feature is the key distinction between device drivers and bus drivers.
- Consumes one or more I/O-related protocols from a controller handle.
- Produces one or more I/O-related protocols on each child handle.

6.2.2 Optional Bus Driver Features

Bus drivers can optionally implement the following features:

- Installs one or more instances of the **EFI_COMPONENT_NAME_PROTOCOL** in the driver's entry point. The implementation of this feature is strongly recommended. It allows a driver to provide human-readable names for the driver and the controllers that the driver is managing. Bus drivers should also provide names or the child handles created by the bus driver. Some platform design guides such as *DIG64* require this feature.
- Installs one or more instances of the **EFI_DRIVER_CONFIGURATION_PROTOCOL** in the driver's entry point. If a driver has any configurable options for the controller or the children, then this protocol is required. Some platform design guides such as *DIG64* require this feature.
- Installs one or more instances of the **EFI_DRIVER_DIAGNOSTICS_PROTOCOL** in the driver's entry point. If a driver wishes to provide diagnostics for the controllers for the children that the driver manages, then this protocol is required. This protocol is the only mechanism that is available to a driver when the driver wants to alert the user of a problem that was detected with a controller. Some platform design guides such as *DIG64* require this feature.
- Provides an **EFI_LOADED_IMAGE_PROTOCOL.Unload()** service, so the driver can be dynamically unloaded. It is recommended that this feature be implemented during driver development, driver debug, and system integration. It is strongly recommended that this service remain in drivers for add-in adapter to help debug interaction issues during system integration.
- Parses the *RemainingDevicePath* parameter that is passed into the **Supported()** and **Start()** services of the **EFI_DRIVER_BINDING_PROTOCOL**.
- Installs an **EFI_DEVICE_PATH_PROTOCOL** on each child handle that is created. This feature is required only if the child handle represents a physical device. If child handle represents a virtual device, then an **EFI_DEVICE_PATH_PROTOCOL** is not required.
- Creates an Exit Boot Services event in the driver's entry point. This feature is required only if the driver is required to place the devices it manages in a specific state just before control is handed to an operating system.
- Creates a Set Virtual Address Map event in the driver's entry point. This feature is required only for a device driver that is also an EFI runtime driver.

6.2.3 Bus Drivers with One Driver Binding Protocol

The driver entry point of a bus driver is very similar to the driver entry point of a device driver. The discussion in section 6.1.3 applies equally well to both bus drivers and device drivers. The differences between bus drivers and device drivers are exposed in the implementations of the **EFI_DRIVER_BINDING_PROTOCOL**. The following sections describe the behaviors of the **Start()** function of the **EFI_DRIVER_BINDING_PROTOCOL** for each type of bus driver.

Table 6-2 below lists the bus drivers from the *EFI 1.10 Sample Implementation* that produce a single instance of the **EFI_DRIVER_BINDING_PROTOCOL**.

Table 6-2. Bus Drivers with One Driver Binding Protocols

DebugPort	Partition	SerialMouse	BiosInt\BiosSnp16
IsaBus	PciBus	Undi	
IsaSerial	ScsiBus	WinNtThunk\SerialIo	

6.2.4 Bus Drivers with Multiple Driver Binding Protocols

The driver entry point of a bus driver is very similar to the driver entry point of a device driver. The discussion in section 6.1.4 applies equally well to both bus drivers and device drivers. The differences between bus drivers and device drivers are exposed in the implementations of the **EFI_DRIVER_BINDING_PROTOCOL**. The following sections describe the behaviors of the **Start()** function of the **EFI_DRIVER_BINDING_PROTOCOL** for each type of bus driver.

There is only one bus driver in the *EFI 1.10 Sample Implementation* that produces multiple instances of the **EFI_DRIVER_BINDING_PROTOCOL**, and this driver is the console splitter driver. This driver multiplexes multiple console output and console input devices into a single virtual console device. It produces instances of the **EFI_DRIVER_BINDING_PROTOCOL** for the following:

- Console output devices
- Standard error device
- Console input device
- Simple pointer devices

This driver is an example of a single feature that can be added or removed from a platform by adding or removing a single component. It could have been implemented as four different drivers, but there were many common functions between the drivers, so it also saved code space to combine these four functions.

6.2.5 Bus Driver Protocol and Child Management

The management of I/O-related protocols by a bus driver is very similar to the management of I/O-related protocol for device drivers that is described in section 6.1.5. A bus driver opens one or more I/O-related protocols on the controller handle for the bus controller, and it creates one or more child handles and installs one or more I/O-related protocols. If the child handle represents a physical device, then a Device Path Protocol must also be installed onto the child handle. The child handle is also required to open the parent I/O protocol with an attribute of **BY_CHILD_CONTROLLER**.

Some types of bus drivers can produce a single child handle each time **Start()** is called, but only if the *RemainingDevicePath* passed into **Start()** represents a valid child device. This distinction means that it may take multiple calls to **Start()** to produce all the child handles. If *RemainingDevicePath* is **NULL**, than all of the remaining child handles will be created at once. When a bus driver opens an I/O-related protocol on the controller handle, it will typically use an open mode of **BY_DRIVER**. However, depending on the type of bus driver, a return code of **EFI_ALREADY_STARTED** from **gBS->OpenProtocol()** may be acceptable. If a device driver gets this return code, then the device driver should not manage the controller handle. If a bus driver gets this return code, then it means that the bus driver has already connected to the controller handle at some point in the past.

Figure 6-6 below shows a simple bus driver that consumes **Protocol A** from a bus controller handle and creates N child handles with a **Device Path Protocol** and **Protocol B**. The **Stop()** function is responsible for destroying the child handles by removing **Protocol B** and the **Device Path Protocol**. Protocol A is first opened **BY_DRIVER** so **Protocol A** cannot be requested by any other drivers. Then, as each child handle is created, the child handle opens **Protocol A** **BY_CHILD_CONTROLLER**. Using this attribute records the parent-child relationship in the handle database, so this information can be extracted if needed. The parent-child links are used by **gBS->DisconnectController()** when a request is made to stop a bus controller.

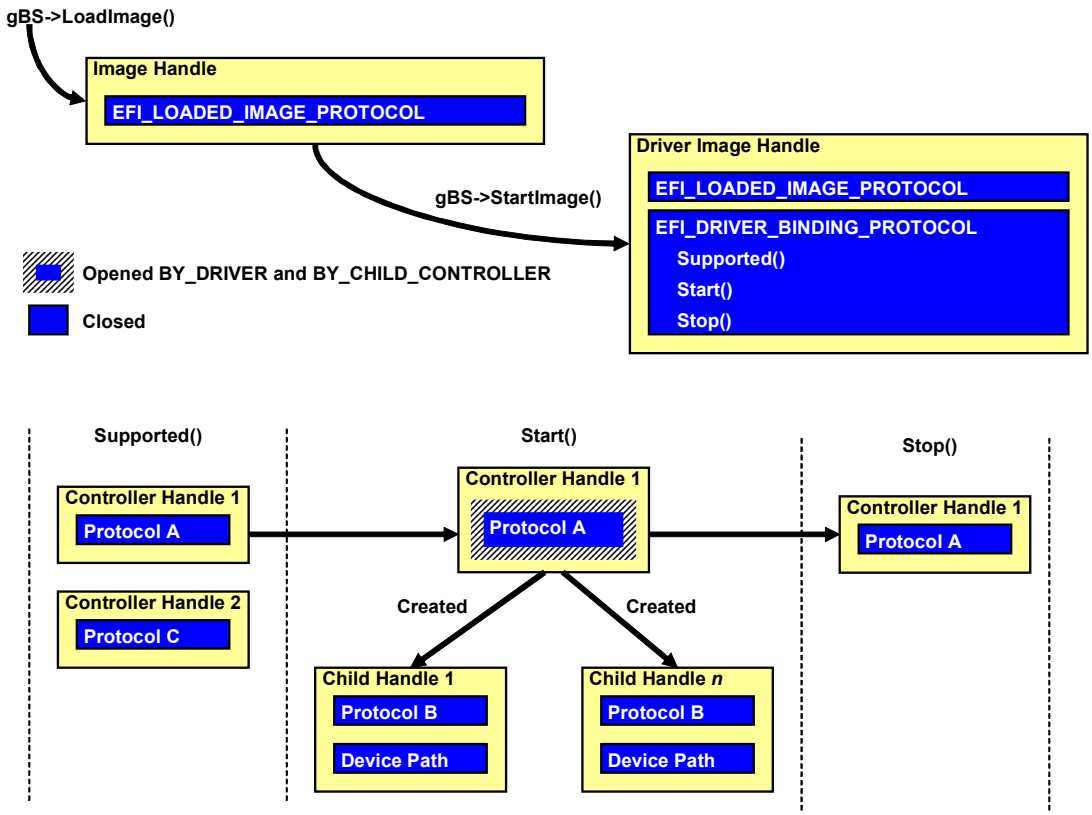

Figure 6-6. Bus Driver Protocol Management

The following sections describe the subtle differences in the child handle creation for each of the bus driver types.

6.2.6 Bus Drivers That Produce One Child in Start()

If the *RemainingDevicePath* parameter passed into **Supported()** and **Start()** is **NULL**, then the bus driver must produce child handles for all the children. If *RemainingDevicePath* is not **NULL**, then the bus driver should parse *RemainingDevicePath* and attempt to produce only the one child device that is described by *RemainingDevicePath*. If the driver does not recognize the device path node in *RemainingDevicePath*, or if the device that is described by the device path node does not match any of the children that are currently attached to the bus controller, then the **Supported()** and **Start()** services should fail. If the *RemainingDevicePath* is recognized and the device path node does match a child device that is attached to the bus controller, then a child handle should be created for that one child device. This step does not make sense for all bus types, because some bus types require the entire bus to be enumerated to produce even a single child. In these cases, the *RemainingDevicePath* should be ignored.

If a bus type has the ability to produce a child handle without enumerating the entire bus, then this ability should be implemented. Implementing this feature will provide faster boot times and is one of the major advantages of the EFI Driver Model. The EFI boot manager may pass the *RemainingDevicePath* of the console device and boot devices to **gBS->ConnectController()**, and **gBS->ConnectController()** will pass this same *RemainingDevicePath* into the **Supported()** and **Start()** services of the **EFI_DRIVER_BINDING_PROTOCOL**. This design allows the minimum number of drivers to be started to boot an operating system. This process can be repeated, so one additional child handle can be produced in each call to **Start()**. Also, a few child handles can be created from the first few calls to **Start()** and then a *RemainingDevicePath* of **NULL** may be passed in, which would required the rest of the child handle to be produced. For example, most SCSI buses do not need to be scanned to create a handle for a SCSI device whose SCSI PUN and SCSI LUN is known ahead of time. By starting only the single hard disk on a SCSI channel that is required to boot an operating system, the scanning of all the other SCSI devices can be eliminated.

Table 6-3 below lists the bus drivers from the *EFI 1.10 Sample Implementation* that can produce a single child handle in the **Start()** service of the **EFI_DRIVER_BINDING_PROTOCOL**.

Table 6-3. Bus Drivers That Produce One Child in Start()

Ide	PciBus	ScsiBus	WinNtThunk\WinNtBusDriver

6.2.7 Bus Drivers That Produce All Children in Start()

If a bus driver is always required to enumerate all of its child devices, then the *RemainingDevicePath* parameter should be ignored in the **Supported()** and **Start()** services of the **EFI_DRIVER_BINDING_PROTOCOL**. All of the child handles should be produced in the first call to **Start()**.

Table 6-4 below lists the bus drivers from the *EFI 1.10 Sample Implementation* that produce all of the child handles in the first call to the **Start()** service of the **EFI_DRIVER_BINDING_PROTOCOL**.

Table 6-4. Bus Drivers That Produce All Children in Start()

Console\ConSplitter	IsaBus	Partition	

6.2.8 Bus Drivers That Produce at Most One Child in Start()

Some bus drivers are for bus controllers that have only a single port, so they have at most one child handle. If *RemainingDevicePath* is **NULL**, then that one child handle should be produced. If *RemainingDevicePath* is not **NULL**, then the *RemainingDevicePath* should be parsed to see if it matches a device path node that the bus driver knows how to produce.

For example, a serial port can have only one device attached to it. This device may be a terminal, a mouse, or a drill press, for example. The driver that consumes the Serial I/O Protocol from a handle must create a child handle with the produced protocol that uses the services of the Serial I/O Protocol.

Table 6-5 below lists the bus drivers from the *EFI 1.10 Sample Implementation* that produce at most one child handle in a call to the **Start()** service of the **EFI_DRIVER_BINDING_PROTOCOL**.

Table 6-5. Bus Drivers That Produce at Most One Child in Start()

Console\Terminal	IsaSerial	Undi	BiosInt\BiosSnp16
DebugPort	SerialMouse	WinNtThunk\SerialIo	

6.2.9 Bus Drivers That Produce No Children in Start()

If a bus controller supports hot-plug devices and the EFI driver wants to support hot-plug events, then no child handles should be produced in **Start()**. Instead, a periodic timer event should be created, and each time the notification function for the periodic timer event is called, the bus driver should check to see if any devices have been hot added or hot removed from the bus. Any devices that were already plugged into the bus when the driver was first started will look like they were just hot added, so the child handles for the devices that were already plugged into the bus will be produced the first time the notification function is executed.

The USB bus driver is the only driver in the *EFI 1.10 Sample Implementation* that produces no children in the **Start()** service of the **EFI_DRIVER_BINDING_PROTOCOL**. This driver is in the **\Efi1.1\Edk\Drivers\Usb\UsbBus** directory.

6.2.10 Bus Drivers That Produce Children with Multiple Parents

Sometimes a bus driver may produce a child handle, and that child handle will actually use the services of multiple parent controllers. This design is useful when a group of parent controllers needs to be multiplexed. The bus driver in this case would manage multiple parent controllers and produce a single child handle. The services produced on that single child handle would make use of the services from each of the parent controllers. Typically, the child device is a virtual device, so a Device Path Protocol would not be installed onto the child handle.

The console splitter bus driver is the only driver in the *EFI 1.10 Sample Implementation* that produces children with multiple parent controllers in the **Start()** service of the **EFI_DRIVER_BINDING_PROTOCOL**. This driver is in the **\Efi1.1\Edk\Drivers\Console\ConSplitter** directory.

6.3 Hybrid Drivers

This driver type has features of both a device driver and a bus driver. The main distinction between a device driver and a bus driver is that a bus driver creates child handles and a device driver does not create any child handles. In addition, a bus driver is allowed only to install produced protocols on the newly created child handles. A hybrid driver does the following:

- Creates new child handles.
- Installs produced protocols on the child handles.
- Installs produced protocols onto the bus controller handle.

A driver for a single-channel PCI SCSI host controller is a hybrid driver. It produces the SCSI Pass Thru Protocol on the controller handle for the PCI SCSI host controller, and it creates child handles for SCSI disk devices and installs the Device Path Protocol and the Block I/O Protocol.

Table 6-6 below lists the hybrid drivers from the *EFI 1.10 Sample Implementation*.

Table 6-6. Hybrid Drivers

Console\Terminal	Ide	Usb\UsbBus	WinNtThunk\WinNtBusDriver

6.3.1 Required Hybrid Driver Features

Hybrid drivers are required to implement the following features:

- Installs one or more instances of the **EFI_DRIVER_BINDING_PROTOCOL** in the driver's entry point.
- Manages one or more controller handles. Even if a driver writer is convinced the driver will manage only a single bus controller, the driver should be designed to manage multiple bus controllers. The overhead for this functionality is low, and it will make the driver more portable.
- Produces child handles. This feature is the key distinction between device drivers and bus drivers.
- Consumes one or more I/O-related protocols from a controller handle.
- Produces one or more I/O-related protocols on the same controller handle.
- Produces one or more I/O-related protocols on each child handle.

6.3.2 Optional Hybrid Driver Features

Hybrid drivers can optionally implement the following features:

- Installs one or more instances of the **EFI_COMPONENT_NAME_PROTOCOL** in the driver's entry point. Implementing this feature is strongly recommended. It allows a driver to provide human-readable names for the driver and the controllers that the driver is managing. Hybrid drivers should also provide names for the child handles created by the hybrid driver. Some platform design guides such as *DIG64* require this feature.
- Installs one or more instances of the **EFI_DRIVER_CONFIGURATION_PROTOCOL** in the driver's entry point. If a driver has any configurable options for the controller or the children, then this protocol is required. Some platform design guides such as *DIG64* require this feature.
- Installs one or more instances of the **EFI_DRIVER_DIAGNOSTICS_PROTOCOL** in the driver's entry point. If a driver wishes to provide diagnostics for the controllers for the children that the driver manages, then this protocol is required. This protocol is the only mechanism that is available to a driver when the driver wants to alert the user of a problem that was detected with a controller. Some platform design guides such as *DIG64* require this feature.
- Provides an **EFI_LOADED_IMAGE_PROTOCOL.Unload()** service, so the driver can be dynamically unloaded. It is recommended that this feature be implemented during driver development, driver debug, and system integration. It is strongly recommended that this service remain in drivers for add-in adapter to help debug interaction issues during system integration.
- Installs an **EFI_DEVICE_PATH_PROTOCOL** on each child handle that is created. This feature is required only if the child handle represents a physical device. If a child handle represents a virtual device, then an **EFI_DEVICE_PATH_PROTOCOL** is not required.
- Creates an Exit Boot Services event in the driver's entry point. This feature is required only if the driver is required to place the devices it manages in a specific state just before control is handed to an operating system.
- Creates a Set Virtual Address Map event in the driver's entry point. This feature is required only for a device driver that is also an EFI runtime driver.

6.4 Service Drivers

A service driver does not manage any devices and it does not produce any instances of the **EFI_DRIVER_BINDING_PROTOCOL**. It is a simple driver that produces one or more protocols on one or more new service handles. These service handles do not have a Device Path Protocol because they do not represent physical devices. The driver entry point returns **EFI_SUCCESS** after the service handles are created and the protocols are installed, which leaves the driver resident in system memory.

Table 6-7 below lists the service drivers from the *EFI 1.10 Sample Implementation*.

Table 6-7. Service Drivers

Bis	DebugSupport	Decompress	Ebc

6.5 Root Bridge Drivers

A root bridge driver does not produce any instances of the **EFI_DRIVER_BINDING_PROTOCOL**. It is responsible for initializing and immediately creating physical controller handles for the root bridge controllers in a platform. The driver must install the Device Path Protocol onto the physical controller handles because the root bridge controllers represent physical devices.

Table 6-8 below lists the root bridge drivers from the *EFI 1.10 Sample Implementation*.

Table 6-8. Root Bridge Drivers

PcatPciRootBridge	WinNtThunk\WinNtPciRootBridge	BiosInt\Disk

6.6 Initializing Drivers

An initializing driver does not create any handles and it does not add any protocols to the handle database. Instead, this type of driver performs some initialization operations and returns an error code so the driver is unloaded from system memory. There are no drivers in the *EFI 1.10 Sample Implementation* of this type.

7
Driver Entry Point

This chapter covers the entry point for the different classes of EFI drivers and their optional features. The most common class of EFI driver is one that follows the EFI Driver Model. This class of driver will be discussed first, followed by the other major types of drivers and the optional features that drivers may choose to implement. Following are the classes of EFI drivers that will be discussed:

- Device, bus, and hybrid drivers that follows the EFI Driver Model
- Initializing driver
- Service driver
- Root bridge driver

Unloadable drivers, Exit Boot Services events, and Set Virtual Address Map events will also be discussed.

The driver entry point is the function that is called when an EFI driver is loaded and started with the **gBS->LoadImage()** and **gBS->StartImage()** services. EFI drivers use the PE/COFF image format that is defined in the *Microsoft Portable Executable and Common Object File Format Specification*. This format supports a single entry point in the code section of the image. The **gBS->StartImage()** service transfers control to the EFI driver at this entry point.

Example 7-1 below shows the entry point to an EFI driver called **Abc**. This example will be expanded upon as each of EFI driver classes and features are discussed. The entry point to an EFI driver is identical to the standard EFI image entry point that is discussed in section 4.1 of the *EFI 1.10 Specification*. The image handle of the EFI driver and a pointer to the EFI System Table are passed into every EFI driver. The image handle allows the EFI driver to discover information about itself, and the pointer to the EFI System Table allows the EFI driver to make EFI Boot Service and EFI Runtime Service calls. The EFI driver can use the EFI Boot Services to access the protocol interfaces that are installed in the handle database, which allows the EFI driver to use the services that are provided through the various protocol interfaces. The driver entry point is preceded by the macro **EFI_DRIVER_ENTRY_POINT()** which declares the driver entry point so source-level debugging is enabled in all build environments.

```
EFI_DRIVER_ENTRY_POINT (AbcDriverEntryPoint)

EFI_STATUS
EFIAPI
AbcDriverEntryPoint (
  IN EFI_HANDLE        ImageHandle,
  IN EFI_SYSTEM_TABLE  *SystemTable
  )
{
}
```

Example 7-1. Generic Entry Point

7.1 EFI Driver Model Driver Entry Point

Drivers that follow the EFI Driver Model are not allowed to touch any hardware in their driver entry point. In fact, these types of drivers do very little in their driver entry point. They are simply required to register interfaces in the handle database by installing protocols. This design allows these types of drivers to be loaded at any point in the system initialization sequence because their driver entry points depend only on a few of the EFI Boot Services. The protocol interfaces that are installed in the driver entry point are used at a later time to initialize, configure, or diagnose the console and boot devices that required to boot an operating system.

All EFI drivers that follow the EFI Driver Model must install one or more instances of the Driver Binding Protocol onto handles in the handle database. The first Driver Binding Protocol is typically installed onto the image handle for the EFI driver. All additional instances of the Driver Binding Protocol should be installed onto new handles.

EFI drivers may optionally support the following:

- Component Name Protocol
- Driver Configuration Protocol
- Driver Diagnostics Protocol

Some platform specifications, such as *DIG64,* require these optional protocols.

EFI Driver Library Services are provided to simplify the driver entry point of an EFI driver. The examples in this section make use of the two EFI Driver Library Services shown in Example 7-2 below. The first library function initializes the EFI Driver Library and installs the Driver Binding Protocol onto the handle specified by `DriverBindingHandle`. `DriverBindingHandle` is typically the same as `ImageHandle`, but if it is **NULL**, then the Driver Binding Protocol is installed onto a newly created handle. The second library function also initializes the EFI Driver Library and it installs all the driver-related protocols onto the handle specified by `DriverBindingHandle`. The optional driver-related protocols are defined to be **OPTIONAL** because they can be **NULL** if a driver does not wish to produce that specific optional protocol. Once again, the `DriverBindingHandle` is typically the same as `ImageHandle`, but if it is **NULL**, then all the driver-related protocols will be installed onto a newly created handle.

```
EFI_STATUS
EfiLibInstallDriverBinding (
  IN EFI_HANDLE                        ImageHandle,
  IN EFI_SYSTEM_TABLE                  *SystemTable,
  IN EFI_DRIVER_BINDING_PROTOCOL       *DriverBinding,
  IN EFI_HANDLE                        DriverBindingHandle
  );

EFI_STATUS
EfiLibInstallAllDriverProtocols (
  IN EFI_HANDLE                        ImageHandle,
  IN EFI_SYSTEM_TABLE                  *SystemTable,
  IN EFI_DRIVER_BINDING_PROTOCOL       *DriverBinding,
  IN EFI_HANDLE                        DriverBindingHandle,
  IN EFI_COMPONENT_NAME_PROTOCOL       *ComponentName,        OPTIONAL
  IN EFI_DRIVER_CONFIGURATION_PROTOCOL *DriverConfiguration,  OPTIONAL
  IN EFI_DRIVER_DIAGNOSTICS_PROTOCOL   *DriverDiagnostics     OPTIONAL
  );
```

Example 7-2. Driver Library Functions

Example 7-3 below shows an example of the entry point to the **Abc** driver that installs the Driver Binding Protocol **gAbcDriverBindingProtocol** and the Component Name Protocol **gAbcComponentName** onto the **Abc** driver's image handle and does not install any of the other optional driver-related protocols. This driver simply returns the status from the EFI Driver Library function **EfiLibInstallAllDriverProtocols()**. See chapter 9 for details on the services and data fields that are produced by the Driver Binding Protocol.

```
EFI_DRIVER_BINDING_PROTOCOL gAbcDriverBinding = {
  AbcDriverBindingSupported,
  AbcDriverBindingStart,
  AbcDriverBindingStop,
  0x10,
  NULL,
  NULL
};

EFI_COMPONENT_NAME_PROTOCOL gAbcComponentName = {
  AbcComponentNameGetDriverName,
  AbcComponentNameGetControllerName,
  "eng"
};

EFI_DRIVER_ENTRY_POINT (AbcDriverEntryPoint)

EFI_STATUS
EFIAPI
AbcDriverEntryPoint (
  IN EFI_HANDLE        ImageHandle,
  IN EFI_SYSTEM_TABLE  *SystemTable
  )
{
  //
  // Initialize a simple EFI driver that follows the EFI Driver Model
  //
  return EfiLibInstallAllDriverProtocols (
           ImageHandle,                // Driver's image handle
           SystemTable,                // EFI System Table Pointer
           &gAbcDriverBinding,         // Required parameters
           ImageHandle,                // Handle for driver-related protocols
           &gAbcComponentName,         // Component Name Procol. May be NULL.
           NULL,                       // Configuration Protocol. May be NULL.
           NULL                        // Diagnostics Protocol. May be NULL.
           );

}
```

Example 7-3. Simple EFI Driver Model Driver Entry Point

Example 7-4 below shows another example of the entry point to the **Abc** driver that installs the Driver Binding Protocol **gAbcDriverBindingProtocol** onto the **Abc** driver's image handle and the optional driver-related protocols. This driver returns the status from the EFI Driver Library function **EfiLibInstallAllDriverProtocols()**. This library function is used if one or more of the optional driver related protocols are being installed.

See chapters 9, 10, 11, and 12 respectively for details on the services and data fields produced by the Driver Binding Protocol, Component Name Protocol, Driver Configuration Protocol, and Driver Diagnostics Protocol.

```
EFI_DRIVER_BINDING_PROTOCOL gAbcDriverBinding = {
  AbcDriverBindingSupported,
  AbcDriverBindingStart,
  AbcDriverBindingStop,
  0x10,
  NULL,
  NULL
};

EFI_COMPONENT_NAME_PROTOCOL gAbcComponentName = {
  AbcComponentNameGetDriverName,
  AbcComponentNameGetControllerName,
  "eng"
};

EFI_DRIVER_CONFIGURATION_PROTOCOL gAbcDriverConfiguration = {
  AbcDriverConfigurationSetOptions,
  AbcDriverConfigurationOptionsValid,
  AbcDriverConfigurationForceDefaults,
  "eng"
};

EFI_DRIVER_DIAGNOSTICS_PROTOCOL gAbcDriverDiagnostics = {
  AbcDriverDiagnosticsRunDiagnostics,
  "eng"
};

EFI_DRIVER_ENTRY_POINT (AbcDriverEntryPoint)

EFI_STATUS
EFIAPI
AbcDriverEntryPoint (
  IN EFI_HANDLE        ImageHandle,
  IN EFI_SYSTEM_TABLE  *SystemTable
  )
{
  //
  // Initialize a complex EFI driver that follows the EFI Driver Model
  //
  return EfiLibInstallAllDriverProtocols (
           ImageHandle,                // Driver's image handle
           SystemTable,                // EFI System Table Pointer
           &gAbcDriverBinding,         // Required parameters
           ImageHandle,                // Handle for driver-related protocols
           &gAbcComponentName,         // Component Name Procol. May be NULL.
           &gAbcDriverConfiguration,   // Configuration Protocol. May be NULL.
           &gAbcDriverDiagnostics      // Diagnostics Protocol. May be NULL.
           );
}
```

Example 7-4. Complex EFI Driver Model Driver Entry Point

7.1.1 Multiple Driver Binding Protocols

If an EFI driver supports more than one parent I/O abstraction, then the driver should produce a Driver Binding Protocol for each of the parent I/O abstractions. For example, an EFI driver could be written to support more than one type of hardware device (for example, ISA and PCI). If code can be shared for the common features of a hardware device, then such a driver should save space and reduce maintenance. The only two drivers in the *EFI 1.10 Sample Implementation* that produce more than one Driver Binding Protocol are the console platform driver and the console splitter driver. These drivers contain multiple Driver Binding Protocols because they depend on multiple console-related parent I/O abstractions.

The first Driver Binding Protocol is typically installed onto the image handle of the EFI driver, and additional Driver Binding Protocols are installed onto new handles. The EFI Driver Library functions that were used in the previous two examples support the creation of new handles by passing in a **NULL** for the fourth argument. Example 7-5 below shows the driver entry point for a driver that produces two instances of the Driver Binding Protocol. The optional driver-related protocols are installed onto the image handle with the first Driver Binding Protocol. See chapters 9, 10, 11, and 12 respectively for details on the services and data fields produced by the Driver Binding Protocol, Component Name Protocol, Driver Configuration Protocol, and Driver Diagnostics Protocol.

```
EFI_DRIVER_BINDING_PROTOCOL gAbcFooDriverBinding = {
  AbcFooDriverBindingSupported,
  AbcFooDriverBindingStart,
  AbcFooDriverBindingStop,
  0x10,
  NULL,
  NULL
};

EFI_DRIVER_BINDING_PROTOCOL gAbcBarDriverBinding = {
  AbcBarDriverBindingSupported,
  AbcBarDriverBindingStart,
  AbcBarDriverBindingStop,
  0x10,
  NULL,
  NULL
};

EFI_COMPONENT_NAME_PROTOCOL gAbcFooComponentName = {
  AbcFooComponentNameGetDriverName,
  AbcFooComponentNameGetControllerName,
  "eng"
};

EFI_DRIVER_CONFIGURATION_PROTOCOL gAbcFooDriverConfiguration = {
  AbcFooDriverConfigurationSetOptions,
  AbcFooDriverConfigurationOptionsValid,
  AbcFooDriverConfigurationForceDefaults,
  "eng"
};

EFI_DRIVER_DIAGNOSTICS_PROTOCOL gAbcFooDriverDiagnostics = {
  AbcFooDriverDiagnosticsRunDiagnostics,
  "eng"
};
```

```
EFI_DRIVER_ENTRY_POINT (AbcDriverEntryPoint)

EFI_STATUS
EFIAPI
AbcDriverEntryPoint (
  IN EFI_HANDLE        ImageHandle,
  IN EFI_SYSTEM_TABLE  *SystemTable
  )
{
  EFI_STATUS  Status;

  //
  // Install first Driver Binding Protocol and the rest of the driver-related
  // protocols onto the driver's image handle
  //
  Status = EfiLibInstallAllDriverProtocols (
             ImageHandle,                  // Driver's image handle
             SystemTable,                  // EFI System Table Pointer
             &gAbcFooDriverBinding,        // Driver Binding Protocol
             ImageHandle,                  // Handle for driver-related
                                           // protocols
             &gAbcFooComponentName,        // Component Name Protocol
             &gAbcFooDriverConfiguration,  // Configuration Protocol
             &gAbcFooDriverDiagnostics     // Diagnostics Protocol
             );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Install second Driver Binding Protocol onto a new handle
  //
  return EfiLibInstallAllDriverProtocols (
           ImageHandle,           // Driver's image handle
           SystemTable,           // EFI System Table Pointer
           gAbcBarDriverBinding,  // Driver Binding Protocol
           NULL,                  // Handle for driver-related protocols
           NULL,                  // Component Name Protocol. May be NULL.
           NULL,                  // Configuration Protocol. May be NULL.
           NULL                   // Diagnostics Protocol. May be NULL.
           );
}
```

Example 7-5. Multiple Driver Binding Protocols

7.1.2 Adding the Unload Feature

Any EFI driver can be made unloadable. This feature is useful for some driver classes, but it may not be useful at all for other driver classes. It does not make sense to add the unload feature to an initializing driver because this class of driver will return an error from the driver entry point, which will force the EFI Image Services to automatically unload the initializing driver.

It usually does not make sense for root bridge drivers or service drivers to add the unload feature. These classes of driver typically produce protocols that are consumed by other EFI drivers to produce basic console functions and boot device abstractions. If a root bridge driver or a service driver is unloaded, then any EFI driver that was using the protocols from those drivers would start to fail. If a root bridge driver or service driver can guarantee that it is not being used by any other EFI components, then they may be unloaded without any adverse side effects.

The **`Unload()`** function can be very helpful. It allows the "unload" command in the EFI Shell to completely remove an EFI driver image from memory and remove all of the driver's handles and protocols from the handle database. If a driver is suspected of causing a bug, it is often helpful to "unload" the driver from the EFI Shell and then proceed to run tests knowing that the driver is no longer present in the platform. In these cases, the **`Unload()`** feature is superior to simply stopping the driver with the **`disconnect`** EFI Shell command. If a driver is just disconnected, then the EFI Shell commands **`connect`** and **`reconnect`** may inadvertently restart the driver.

The unload feature is also very helpful when testing and developing new versions of the driver. The old version can be completely unloaded (removed from the system) and new versions of the driver, which may have the same version number, can safely be installed in the system without concerns that the older version of the driver may get invoked during the next connect or reconnect operation.

Because **`Unload()`** completely removes the driver from system memory, it might not be possible to load it back into the system. For example, if the driver is stored in system firmware or in a PCI option ROM, no method may be available for reloading the driver without completely rebooting the system.

The **`Unload()`** service is one of the fields in the **`EFI_LOADED_IMAGE_PROTOCOL`**. This protocol is automatically created and installed when an EFI image is loaded with the EFI Boot Service **`LoadImage()`**. When the **`EFI_LOADED_IMAGE_PROTOCOL`** is created by **`LoadImage()`**, the **`Unload()`** service is initialized to **`NULL`**. It is the driver entry point's responsibility to register the **`Unload()`** function in the **`EFI_LOADED_IMAGE_PROTOCOL`**.

It is recommended that EFI drivers that follow the EFI Driver Model add the unload feature. This feature is very useful during driver development, driver debug, and system integration. It is strongly recommended that this service remain in drivers for add-in adapters to help debug interaction issues during system integration.

Example 7-6 below shows the same driver entry point from Example 7-3 with the unload feature added. Only a template for the **`Unload()`** function is shown in this example because the implementation of this service will vary from driver to driver. The **`Unload()`** service is responsible for cleaning up everything that the driver has done since it was initialized. This responsibility means that the **`Unload()`** service should do the following:

- Free any resources that were allocated.
- Remove any protocols that were added.

- Destroy any handles that were created.

If for some reason the **Unload()** service does not want to unload the driver at the time the **Unload()** service is called, it can return an error and the driver will not be unloaded. The only way that a driver is actually unloaded is if the **Unload()** service has been registered in the **EFI_LOADED_IMAGE_PROTOCOL** and the **Unload()** service returns **EFI_SUCCESS**.

```
EFI_DRIVER_BINDING_PROTOCOL gAbcDriverBinding = {
  AbcDriverBindingSupported,
  AbcDriverBindingStart,
  AbcDriverBindingStop,
  0x10,
  NULL,
  NULL
};

EFI_COMPONENT_NAME_PROTOCOL gAbcComponentName = {
  AbcComponentNameGetDriverName,
  AbcComponentNameGetControllerName,
  "eng"
};

EFI_STATUS
EFIAPI
AbcUnload (
  IN EFI_HANDLE  ImageHandle
  )
{
  return EFI_SUCCESS;
}

EFI_DRIVER_ENTRY_POINT (AbcDriverEntryPoint)

EFI_STATUS
EFIAPI
AbcDriverEntryPoint (
  IN EFI_HANDLE        ImageHandle,
  IN EFI_SYSTEM_TABLE  *SystemTable
  )
{
  EFI_STATUS                  Status;
  EFI_LOADED_IMAGE_PROTOCOL  *LoadedImage;

  //
  // Retrieve the Loaded Image Protocol from image handle
  //
  Status = gBS->OpenProtocol (
                 ImageHandle,
                 &gEfiLoadedImageProtocolGuid,
                 (VOID **)&LoadedImage,
                 ImageHandle,
                 ImageHandle,
                 EFI_OPEN_PROTOCOL_GET_PROTOCOL
                 );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
```

```
  // Fill in the Unload() service of the Loaded Image Protocol
  //
  LoadedImage->Unload = AbcUnload;

  //
  // Initialize a simple EFI driver that follows the EFI Driver Model
  //
  return EfiLibInstallAllDriverProtocols (
           ImageHandle,                // Driver's image handle
           SystemTable,                // EFI System Table Pointer
           &gAbcDriverBinding,         // Required parameters
           ImageHandle,                // Handle for driver related protocols
           &gAbcComponentName,         // Component Name Procol. May be NULL.
           NULL,                       // Configuration Protocol. May be NULL.
           NULL                        // Diagnostics Protocol. May be NULL.
           );
}
```

Example 7-6. Add the Unload Feature

Example 7-7 below shows one possible implementation of the **`Unload()`** function for an EFI driver that follows the EFI Driver Model. It finds all the devices it is managing and disconnects the driver from those devices. Then, the protocol interfaces that were installed in the driver entry point must be removed. The example shown here matches the driver entry point from Example 7-4. There are many possible algorithms that can be implemented in the unload service. A driver may choose to be unloadable if and only if it is not managing any devices at all. A driver may also choose to internally keep track of the devices it is managing, so it can selectively disconnect itself from those devices when it is unloaded.

```
EFI_STATUS
EFIAPI
AbcUnload (
  IN EFI_HANDLE  ImageHandle
  )
{
  EFI_STATUS  Status;
  EFI_HANDLE  *DeviceHandleBuffer;
  UINTN       DeviceHandleCount;
  UINTN       Index;

  //
  // Get the list of all the handles in the handle database.
  // If there is an error getting the list, then the unload operation fails.
  //
  Status = gBS->LocateHandleBuffer (
                AllHandles,
                NULL,
                NULL,
                &DeviceHandleCount,
                &DeviceHandleBuffer
                );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Disconnect the driver specified by ImageHandle from all the devices in the
  // handle database.
```

```
  //
  for (Index = 0; Index < DeviceHandleCount; Index++) {
    Status = gBS->DisconnectController (
                    DeviceHandleBuffer[Index],
                    ImageHandle,
                    NULL
                    );
  }

  //
  // Free the buffer containing the list of handles from the handle database
  //
  if (DeviceHandleBuffer != NULL) {
    gBS->FreePool(DeviceHandleBuffer);
  }

  //
  // Uninstall all the protocols that were installed in the driver entry point
  //
  Status = gBS->UninstallMultipleProtocolInterfaces (
                  ImageHandle,
                  &gEfiDriverBindingProtocolGuid, &gAbcDriverBinding,
                  &gEfiComponentNameProtocolGuid, &gAbcComponentName,
                  &gEfiDriverConfigurationProtocolGuid, &gAbcDriverConfiguration,
                  &gEfiDriverDiagnosticsProtocolGuid, &gAbcDriverDiagnostics
                  NULL
                  );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Do any additional cleanup that is required for this driver
  //

  return EFI_SUCCESS;
}
```

Example 7-7. Unload Function

7.1.3 Adding the Exit Boot Services Feature

Some EFI drivers may need to put their devices into a known state prior to booting an operating system. This case is considered to be very rare because the OS-present drivers should not depend on an EFI driver running at all. Not depending on a running EFI driver means that an OS-present driver should be able to handle the following:

- A device in its power-on reset state
- A device that was recently hot added while the OS is running
- A device that was managed by an EFI driver up to the point where the OS was booted
- A device that was managed for a short period of time by an EFI driver

None of the drivers in the *EFI Sample Implementation* use the feature described here. It is documented here to show what is possible if this feature is ever required.

In the rare case where an EFI driver is required to place a device in a known state before booting an operating system, the driver can use a special event type called an *Exit Boot Services event.* This

event is signaled when an OS loader or OS kernel calls the EFI Boot Service **gBS->ExitBootServices()**. This call is the point in time where the system firmware still owns the platform, but the system firmware is just about to transfer system ownership to the operating system. In this transition time, no modifications to the EFI memory map are allowed (see section 5.1 of the *EFI 1.10 Specification*). This requirement means that the notification function for an Exit Boot Services event is not allowed to directly or indirectly allocate or free and memory through the EFI Memory Services.

Example 7-8 below shows the same example from Example 7-3, but an Exit Boot Services event is also created. The template for the notification function for the Exit Boot Services event is also shown. This notification function will typically contain code to find the list of device handles that the driver is currently managing, and it will then perform operations on those handles to make sure they are in the proper OS handoff state. Remember that no memory allocation or free operations can be performed from this notification function.

```
EFI_DRIVER_BINDING_PROTOCOL gAbcDriverBinding = {
  AbcDriverBindingSupported,
  AbcDriverBindingStart,
  AbcDriverBindingStop,
  0x10,
  NULL,
  NULL
};

EFI_COMPONENT_NAME_PROTOCOL gAbcComponentName = {
  AbcComponentNameGetDriverName,
  AbcComponentNameGetControllerName,
  "eng"
};

VOID
EFIAPI
AbcNotifyExitBootServices (
  IN EFI_EVENT  Event,
  IN VOID       *Context
  )
{
  //
  // Put driver-specific actions here.
  // No EFI Memory Service may be used directly or indirectly.
  //
}

EFI_DRIVER_ENTRY_POINT (AbcDriverEntryPoint)

EFI_STATUS
EFIAPI
AbcDriverEntryPoint (
  IN EFI_HANDLE        ImageHandle,
  IN EFI_SYSTEM_TABLE  *SystemTable
  )
{
  EFI_STATUS  Status;
  EFI_EVENT   *ExitBootServicesEvent;

  //
  // Create an Exit Boot Services event.
```

```
  //
  Status = gBS->CreateEvent (
                 EFI_EVENT_SIGNAL_EXIT_BOOT_SERVICES,
                 EFI_TPL_NOTIFY,
                 AbcNotifyExitBootServices,
                 NULL,
                 &ExitBootServicesEvent
                 );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Initialize a simple EFI driver that follows the EFI Driver Model
  //
  return EfiLibInstallAllDriverProtocols (
           ImageHandle,                // Driver's image handle
           SystemTable,                // EFI System Table Pointer
           &gAbcDriverBinding,         // Required parameters
           ImageHandle,                // Handle for driver-related protocols
           &gAbcComponentName,         // Component Name Procol. May be NULL.
           NULL,                       // Configuration Protocol. May be NULL.
           NULL                        // Diagnostics Protocol. May be NULL.
           );

}
```

Example 7-8. Adding the Exit Boot Services Feature

If an EFI driver supports both the unload feature and the Exit Boot Services feature, then the **`Unload()`** function must destroy the Exit Boot Services event by calling **`gBS->CloseEvent()`**. In this case, the Exit Boot Services event would likely be declared as a global variable, so the event could easily be destroyed in the **`Unload()`** function. Example 7-9 below shows a driver that supports both the unload service and an Exit Boot Services event.

```
EFI_EVENT  *gExitBootServicesEvent;

VOID
EFIAPI
AbcNotifyExitBootServices (
  IN EFI_EVENT  Event,
  IN VOID       *Context
  )
{
  //
  // Put driver-specific actions here.
  // No EFI Memory Service may be used directly or indirectly.
  //
}

EFI_STATUS
EFIAPI
AbcUnload (
  IN EFI_HANDLE  ImageHandle
  )
{
  gBS->CloseEvent (gExitBootServicesEvent);
}
```

```
EFI_DRIVER_ENTRY_POINT (AbcDriverEntryPoint)

EFI_STATUS
EFIAPI
AbcDriverEntryPoint (
  IN EFI_HANDLE        ImageHandle,
  IN EFI_SYSTEM_TABLE  *SystemTable
  )
{
  EFI_STATUS                 Status;
  EFI_LOADED_IMAGE_PROTOCOL  *LoadedImage;

  //
  // Initialize the EFI Driver Library
  //
  EfiInitializeDriverLib (ImageHandle, SystemTable);

  //
  // Create an Exit Boot Services event.
  //
  Status = gBS->CreateEvent (
                 EFI_EVENT_SIGNAL_EXIT_BOOT_SERVICES,
                 EFI_TPL_NOTIFY,
                 AbcNotifyExitBootServices,
                 NULL,
                 &gExitBootServicesEvent
                 );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Retrieve the Loaded Image Protocol from image handle
  //
  Status = gBS->OpenProtocol (
                 ImageHandle,
                 &gEfiLoadedImageProtocolGuid,
                 (VOID **)&LoadedImage,
                 ImageHandle,
                 ImageHandle,
                 EFI_OPEN_PROTOCOL_GET_PROTOCOL
                 );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Fill in the Unload() service of the Loaded Image Protocol
  //
  LoadedImage->Unload = AbcUnload;

  return EFI_SUCCESS;
}
```

Example 7-9. Add the Unload and Exit Boot Services Event Feature

7.2 Initializing Driver Entry Point

Example 7-10 below shows a template for an initializing driver called **Abc**. This driver initializes one or more components in the platform and exits. It does not produce any services that are required after the entry point has been executed. This type of driver will return an error from the entry point, so the driver will be unloaded by the EFI Image Services. An initializing driver will never produce an **Unload()** service because they are always unloaded after their entry point is executed. This type of driver is not used by IHVs. Instead, it is typically used by OEMs and IBVs to initialize the state of a hardware component in the platform such as the processor or chipset components.

```
EFI_DRIVER_ENTRY_POINT (AbcDriverEntryPoint)

EFI_STATUS
EFIAPI
AbcDriverEntryPoint (
  IN EFI_HANDLE        ImageHandle,
  IN EFI_SYSTEM_TABLE  *SystemTable
  )
{
  //
  // Initialize the EFI Driver Library
  //
  EfiInitializeDriverLib (ImageHandle, SystemTable);

  //
  // Perform some platform initialization operations here
  //

  return EFI_LOAD_ERROR;
}
```

Example 7-10. Initializing Driver Entry Point

7.3 Service Driver Entry Point

A service driver produces one or more protocol interfaces on the driver's image handle or on newly created handles. Example 7-11 below shows the Decompress Protocol being installed onto the driver's image handle. A service driver may produce an **Unload()** service, and that service would be required to uninstall the protocols that were installed in the driver's entry point. The **Unload()** service can be a dangerous operation because there is no way for the service driver to know if the protocols that it installed are being used by other EFI components. If the service driver is unloaded and other EFI components are still using the protocols that were produced by the unloaded driver, then the system will likely fail.

```
EFI_DECOMPRESS_PROTOCOL gAbcDecompress = {
  GetInfo,
  Decompress
};

EFI_DRIVER_ENTRY_POINT (AbcDriverEntryPoint)

EFI_STATUS
EFIAPI
```

```
AbcDriverEntryPoint (
  IN EFI_HANDLE        ImageHandle,
  IN EFI_SYSTEM_TABLE  *SystemTable
  )
{
  //
  // Initialize the EFI Driver Library
  //
  EfiInitializeDriverLib (ImageHandle, SystemTable);

  //
  // Install the Decompress Protocol onto the driver's image handle
  //
  return gBS->InstallMultipleProtocolInterfaces (
                &ImageHandle,
                &gEfiDecompressProtocolGuid, &gAbcDecompress,
                NULL
                );
}
```

Example 7-11. Service Driver Entry Point – Image Handle

A service driver may also install its protocol interfaces onto one or more new handles in the handle database. Example 7-12 below shows a template for a service driver called **Abc** that produces the Decompress Protocol on a new handle.

```
EFI_DECOMPRESS_PROTOCOL gAbcDecompress = {
  GetInfo,
  Decompress
};

EFI_DRIVER_ENTRY_POINT (AbcDriverEntryPoint)

EFI_STATUS
EFIAPI
AbcDriverEntryPoint (
  IN EFI_HANDLE        ImageHandle,
  IN EFI_SYSTEM_TABLE  *SystemTable
  )
{
  EFI_HANDLE  NewHandle;

  //
  // Initialize the EFI Driver Library
  //
  EfiInitializeDriverLib (ImageHandle, SystemTable);

  //
  // Install the Decompress Protocol onto the a new handle
  //
  NewHandle = NULL;
  return gBS->InstallMultipleProtocolInterfaces (
                &NewHandle,
                &gEfiDecompressProtocolGuid, &gAbcDecompress,
                NULL
                );
}
```

Example 7-12. Service Driver Entry Point – New Handle

7.4 Root Bridge Driver Entry Point

Root bridge drivers produce handles and software abstractions for the bus types that are directly produced by a core chipset. The PCI Root Bridge I/O Protocol is the software abstraction for root bridges that is defined in the *EFI 1.10 Specification.* EFI drivers that produce this protocol do not follow the EFI Driver Model. Instead, they initialize hardware and directly produce the handles and protocols in the driver entry point. Root bridge drivers are slightly different from service drivers in the following ways:

- Root bridge drivers always create new handles.
- The root bridge driver is responsible for installing both of the following:
  - The software abstraction, such as the PCI Root Bridge I/O Protocol
  - The Device Path Protocol that describes the programmatic path to the root bridge device

Example 7-13 below shows a template for a root bridge driver that produces one handle for a system with a single PCI root bridge. A Device Path Protocol with an ACPI device path node and the PCI Root Bridge I/O Protocol are installed onto a newly created handle. The ACPI device path node for the PCI root bridge must match the description of the PCI root bridge in the ACPI table for the platform. In this example, the Device Path Protocol and PCI Root Bridge I/O Protocol are declared as global variables. An actual driver may need to keep track of additional private data fields to properly manage a PCI root bridge. In that case, global variables would not be used. Instead, a constructor function would allocate and initialize the public and private data fields. This concept is described in more detail in chapter 8.

```
typedef struct {
  ACPI_HID_DEVICE_PATH      AcpiDevicePath;
  EFI_DEVICE_PATH_PROTOCOL  EndDevicePath;
} EFI_PCI_ROOT_BRIDGE_DEVICE_PATH;

EFI_PCI_ROOT_BRIDGE_DEVICE_PATH  gAbcPciRootBridgeDevicePath = {
  {
    ACPI_DEVICE_PATH,                                 // Type
    ACPI_DP,                                          // Subtype
    (UINT8) (sizeof(ACPI_HID_DEVICE_PATH)),           // Length (lower 8 bits)
    (UINT8) ((sizeof(ACPI_HID_DEVICE_PATH)) >> 8),    // Length (upper 8 bits)
    EISA_PNP_ID(0x0A03),                              // HID
    0                                                 // UID
  },
  {
    END_DEVICE_PATH_TYPE,                             // Type
    END_ENTIRE_DEVICE_PATH_SUBTYPE,                   // Subtype
    END_DEVICE_PATH_LENGTH,                           // Length (lower 8 bits)
    0                                                 // Length (upper 8 bits)
  }
};

EFI_PCI_ROOT_BRIDGE_IO_PROTOCOL gAbcPciRootBridgeIo = {
  NULL,                                  // ParentHandle
  AbcPciRootBridgeIoPollMem,             // PollMem()
  AbcPciRootBridgeIoPollIo,              // PolIo()
  {
    AbcPciRootBridgeIoMemRead,           // Mem.Read()
    AbcPciRootBridgeIoMemWrite           // Mem.Write()
  },
  {
```

```
    AbcPciRootBridgeIoIoRead,          // Io.Read()
    AbcPciRootBridgeIoIoWrite,         // Io.Write()
  },
  {
    AbcPciRootBridgeIoPciRead,         // Pci.Read()
    AbcPciRootBridgeIoPciWrite,        // Pci.Write()
  },
  AbcPciRootBridgeIoCopyMem,           // CopyMem()
  AbcPciRootBridgeIoMap,               // Map()
  AbcPciRootBridgeIoUnmap,             // Unmap()
  AbcPciRootBridgeIoAllocateBuffer,    // AllocateBuffer()
  AbcPciRootBridgeIoFreeBuffer,        // FreeBuffer()
  AbcPciRootBridgeIoFlush,             // Flush()
  AbcPciRootBridgeIoGetAttributes,     // GetAttributes()
  AbcPciRootBridgeIoSetAttributes,     // SetAttributes()
  AbcPciRootBridgeIoConfiguration,     // Configuration()
  0                                    // SegmentNumber
};

EFI_DRIVER_ENTRY_POINT (AbcDriverEntryPoint)

EFI_STATUS
EFIAPI
AbcDriverEntryPoint (
  IN EFI_HANDLE        ImageHandle,
  IN EFI_SYSTEM_TABLE  *SystemTable
  )
{
  EFI_STATUS  Status;
  EFI_HANDLE  NewHandle;

  //
  // Initialize the EFI Driver Library
  //
  EfiInitializeDriverLib (ImageHandle, SystemTable);

  //
  // Perform root bridge initialization operations here
  //

  //
  // Install the Device Path Protocol and PCI Root Bridge I/O Protocol onto
  // a new handle.
  //
  NewHandle = NULL;
  Status = gBS->InstallMultipleProtocolInterfaces (
                 &NewHandle,
                 &gEfiDevicePathProtocolGuid,      &gAbcPciRootBridgeDevicePath,
                 &gEfiPciRootBridgeIoProtocolGuid, &gAbcPciRootBridgeIo,
                 NULL
                 );
  return Status;
}
```

Example 7-13. Single PCI Root Bridge Driver Entry Point

Example 7-14 below shows a template for a root bridge driver that produces four handles for a system with four PCI root bridges. A Device Path Protocol with an ACPI device path node and the PCI Root Bridge I/O Protocol are installed onto each of the newly created handles. The ACPI device path nodes for each of the PCI root bridges must match the description of the PCI root bridges in the ACPI tables for the platform. In this example, the _UID field for the root bridges has the values of 0, 1, 2, and 3. However, there is no requirement that the _UID field starts at 0 or that they are contiguous. The only requirement is that the _UID field for each root bridge matches the _UID field in the ACPI table that describes the same root bridge controller. The Device Path Protocol and PCI Root Bridge I/O Protocol are declared as global variables, and copies of those global variables are made for each PCI root bridge. An actual driver may need to keep track of additional private data fields to properly manage each PCI root bridge. In that case, global variables would not be used. Instead, a constructor function would allocate and initialize the public and private data fields. This concept is described in more detail in chapter 8.

```
#define HWP_EISA_ID_CONST  0x22F0
#define COMPRESSED_ASCII (C1, C2, C3) ((((C1) - 'A') & 0x1f) << 10) | \
                                      ((((C2) - 'A') & 0x1f) << 5 ) | \
                                      ((((C3) - 'A') & 0x1f)

typedef struct {
  ACPI_ EXTENDED_HID_DEVICE_PATH  AcpiDevicePath;
  EFI_DEVICE_PATH_PROTOCOL        EndDevicePath;
} EFI_PCI_ROOT_BRIDGE_DEVICE_PATH;

EFI_PCI_ROOT_BRIDGE_DEVICE_PATH  gAbcPciRootBridgeDevicePath = {
  {
    ACPI_DEVICE_PATH,                                // Type
    ACPI_DP,                                         // Subtype
    (UINT8) (sizeof(ACPI_HID_DEVICE_PATH)),          // Length (lower 8 bits)
    (UINT8) ((sizeof(ACPI_HID_DEVICE_PATH)) >> 8),   // Length (upper 8 bits)
    EISA_ID (COMPRESSED_ASCII ('H','W','P'), 2),     // HID
    0,                                               // UID
    EISA_PNP_ID (0x0A03)                             // CID
  },
  {
    END_DEVICE_PATH_TYPE,                            // Type
    END_ENTIRE_DEVICE_PATH_SUBTYPE,                  // Subtype
    END_DEVICE_PATH_LENGTH,                          // Length (lower 8 bits)
    0                                                // Length (upper 8 bits)
  }
};

EFI_PCI_ROOT_BRIDGE_IO_PROTOCOL gAbcPciRootBridgeIo = {
  NULL,                                 // ParentHandle
  AbcPciRootBridgeIoPollMem,            // PollMem()
  AbcPciRootBridgeIoPollIo,             // PolIo()
  {
    AbcPciRootBridgeIoMemRead,          // Mem.Read()
    AbcPciRootBridgeIoMemWrite          // Mem.Write()
  },
  {
    AbcPciRootBridgeIoIoRead,           // Io.Read()
    AbcPciRootBridgeIoIoWrite,          // Io.Write()
  },
  {
    AbcPciRootBridgeIoPciRead,          // Pci.Read()
    AbcPciRootBridgeIoPciWrite,         // Pci.Write()
```

```
  },
  AbcPciRootBridgeIoCopyMem,           // CopyMem()
  AbcPciRootBridgeIoMap,               // Map()
  AbcPciRootBridgeIoUnmap,             // Unmap()
  AbcPciRootBridgeIoAllocateBuffer,    // AllocateBuffer()
  AbcPciRootBridgeIoFreeBuffer,        // FreeBuffer()
  AbcPciRootBridgeIoFlush,             // Flush()
  AbcPciRootBridgeIoGetAttributes,     // GetAttributes()
  AbcPciRootBridgeIoSetAttributes,     // SetAttributes()
  AbcPciRootBridgeIoConfiguration,     // Configuration()
  0                                    // SegmentNumber
};

EFI_DRIVER_ENTRY_POINT (AbcDriverEntryPoint)

EFI_STATUS
EFIAPI
AbcDriverEntryPoint (
  IN EFI_HANDLE        ImageHandle,
  IN EFI_SYSTEM_TABLE  *SystemTable
  )
{
  EFI_STATUS                       Status;
  UINTN                            Index;
  EFI_HANDLE                       NewHandle;
  EFI_PCI_ROOT_BRIDGE_DEVICE_PATH  *DevicePath;
  EFI_PCI_ROOT_BRIDGE_IO_PROTOCOL  *PciRootBridgeIo;

  //
  // Initialize the EFI Driver Library
  //
  EfiInitializeDriverLib (ImageHandle, SystemTable);

  //
  // Perform root bridge initialization operations here
  //

  //
  // Install the Device Path Protocol and PCI Root Bridge I/O Protocol onto
  // a new handle.
  //
  for (Index = 0; Index < 4 ; Index++) {

    //
    // Make a copy of the device path and update the UID field of the ACPI
    // device path node
    //
    DevicePath = EfiLibAllocatePool (sizeof (EFI_PCI_ROOT_BRIDGE_DEVICE_PATH));
    if (DevicePath == NULL) {
      return EFI_OUT_OF_RESOURCES;
    }
    gBS->CopyMem (
           DevicePath,
           &gAbcPciRootBridgeDevicePath,
           sizeof (EFI_PCI_ROOT_BRIDGE_DEVICE_PATH)
           );
    DevicePath->AcpiDevicePath.UID = Index;

    //
```

```
    // Make a copy of the PCI Root Bridge I/O Protocol
    //
    PciRootBridgeIo = EfiLibAllocatePool (
                        sizeof (EFI_PCI_ROOT_BRIDGE_IO_PROTOCOL)
                        );
    if (PciRootBridgeIo == NULL) {
      return EFI_OUT_OF_RESOURCES;
    }
    gBS->CopyMem (
           PciRootBridgeIo,
           &gAbcPciRootBridgeIo,
           sizeof (EFI_PCI_ROOT_BRIDGE_IO_PROTOCOL)
           );

    //
    // Install the Device Path Protocol and the PCI Root Bridge I/O Protocol
    // onto a new handle.
    //
    NewHandle = NULL;
    Status = gBS->InstallMultipleProtocolInterfaces (
                    &NewHandle,
                    &gEfiDevicePathProtocolGuid,  DevicePath,
                    &gEfiPciRootBridgeIoProtocolGuid, PciRootBridgeIo,
                    NULL
                    );
    if (EFI_ERROR (Status)) {
      return Status;
    }
  }
  return EFI_SUCCESS;
}
```

Example 7-14. Multiple PCI Root Bridge Driver Entry Point

7.5 Runtime Drivers

Any EFI driver can be modified to be a runtime driver. This ability applies to the following:

- EFI drivers that follow the EFI Driver Model
- Initializing drivers
- Service drivers
- Root bridge drivers

However, it does not make sense to add the runtime feature to an initializing driver, because the initializing driver is unloaded after its entry point has been executed. The best example of a runtime driver that follows the EFI Driver Model is an UNDI driver that provides services for a network interface controller (NIC).

A runtime driver is a driver that is required to produce services after **`gBS->ExitBootServices()`** has been called. Drivers of this type are much more difficult to implement and validate because they are required to execute in both the preboot environment where the system firmware owns the platform and while an OS is running where the OS owns the platform. The OS may choose to switch all runtime services from physical mode addressing to virtual mode addressing. The driver cannot know which type of OS is going to be booted, so the

runtime driver must always be able to switch from physical addressing to virtual addressing if **gRT->SetVirtualAddressMap()** is called by an OS loader or an OS kernel. In addition, because all memory regions that are marked as boot services memory in the EFI memory map are converted to available memory when an OS is booted, a runtime driver must allocate all of its memory buffers from runtime memory.

A runtime driver will typically create the following two special events to help with these issues:

- Exit Boot Services event
- Set Virtual Address Map event

The Exit Boot Services event is signaled when the OS loader or OS kernel calls **gBS->ExitBootServices()**. After this point, the EFI driver is not allowed to use any of the EFI Boot Services. The EFI Runtime Services and services from other runtime drivers are still available. The Set Virtual Address Map event is signaled when the OS loader or OS kernel calls **gRT->SetVirtualAddresMap()**. If this event is signaled, then the OS loader or OS kernel is requesting that all runtime components be converted from their physical address mapping to the virtual address mappings that are passed to **gRT->SetVirtualAddressMap()**. The EFI firmware does most of the work here by relocating all the EFI images from their physically addressed code and data segments to their virtually addressed code and data segments. However, the EFI firmware cannot know what memory buffers a runtime component has allocated and what pointer values stored within those runtime memory buffers need to be converted from their physical addresses to their virtual addresses. The notification function for the Set Virtual Address Map event is required to use the **gRT->ConvertPointer()** service to convert all pointers in private data structures from their physical address to their virtual addresses. This code is complex and difficult to get correct because no tools are available at this time to help know when all the pointers have been converted. The only symptom that is seen when it is not done correctly is that the OS will crash in the middle of a call to a service produced by a runtime driver.

Example 7-15 below shows the driver entry point for a runtime driver that creates an Exit Boot Services event and a Set Virtual Address Map event. The notification function for the Exit Boot Services event sets a global variable to **TRUE**, so the code in other functions can know if the EFI Boot Services are available. This global variable is initialized to **FALSE** in its declaration. The notification function for the Set Virtual Address Map event converts one global pointer from a physical address to a virtual address as an example. A real driver might have many more pointers to convert. In general, a runtime driver should be designed to reduce or eliminate pointers that need to be converted to minimize the likelihood of missing a pointer conversion.

```
VOID     *gGlobalPointer;
BOOLEAN  gAtRuntime = FALSE;

VOID
EFIAPI
AbcNotifySetVirtualAddressMap (
  IN EFI_EVENT  Event,
  IN VOID       *Context
  )
{
  gRT->ConvertPointer (
         EFI_OPTIONAL_POINTER,
         (VOID **)&gGlobalPointer
         );
}
```

```
VOID
EFIAPI
AbcNotifyExitBootServices (
  IN EFI_EVENT  Event,
  IN VOID       *Context
  )
{
  gAtRuntime = TRUE;
}

EFI_DRIVER_ENTRY_POINT (AbcDriverEntryPoint)

EFI_STATUS
EFIAPI
AbcDriverEntryPoint (
  IN EFI_HANDLE        ImageHandle,
  IN EFI_SYSTEM_TABLE  *SystemTable
  )
{
  EFI_STATUS  Status;
  EFI_EVENT   *ExitBootServicesEvent;
  EFI_EVENT   *SetVirtualAddressMapEvent;

  //
  // Create an Exit Boot Services event.
  //
  Status = gBS->CreateEvent (
                 EFI_EVENT_SIGNAL_EXIT_BOOT_SERVICES,
                 EFI_TPL_NOTIFY,
                 AbcNotifyExitBootServices,
                 NULL,
                 &ExitBootServicesEvent
                 );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Create a Set Virtual Address Map event.
  //
  Status = gBS->CreateEvent (
                 EFI_EVENT_SIGNAL_SET_VIRTUAL_ADDRESS_MAP,
                 EFI_TPL_NOTIFY,
                 AbcNotifySetVirtualAddressMap,
                 NULL,
                 &SetVirtualAddressMapEvent
                 );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Perform additional driver initialization here
  //

  return EFI_SUCCESS;
}
```

Example 7-15. Runtime Driver Entry Point

If a runtime driver also supports the unload feature, then the **Unload()** function must destroy the Exit Boot Services event and the Set Virtual Address Map event by calling **gBS->CloseEvent()**. In this case, these events would likely be declared as global variables, so the events could easily be destroyed in the **Unload()** function. Example 7-16 below shows an unloadable runtime driver.

```
EFI_EVENT  *gExitBootServicesEvent;
EFI_EVENT  *gSetVirtualAddressMapEvent;
VOID       *gGlobalPointer;
BOOLEAN    gAtRuntime = FALSE;

VOID
EFIAPI
AbcNotifySetVirtualAddressMap (
  IN EFI_EVENT  Event,
  IN VOID       *Context
  )
{
  gRT->ConvertPointer (
         EFI_OPTIONAL_POINTER,
         (VOID **)&gGlobalPointer
         );
}

VOID
EFIAPI
AbcNotifyExitBootServices (
  IN EFI_EVENT  Event,
  IN VOID       *Context
  )
{
  gAtRuntime = TRUE;
}

EFI_STATUS
EFIAPI
AbcUnload (
  IN EFI_HANDLE  ImageHandle
  )
{
  gBS->CloseEvent (gExitBootServicesEvent);
  gBS->CloseEvent (gSetVirtualAddressMapEvent);
}

EFI_DRIVER_ENTRY_POINT (AbcDriverEntryPoint)

EFI_STATUS
EFIAPI
AbcDriverEntryPoint (
  IN EFI_HANDLE        ImageHandle,
  IN EFI_SYSTEM_TABLE  *SystemTable
  )
{
  EFI_STATUS                 Status;
  EFI_LOADED_IMAGE_PROTOCOL  *LoadedImage;

  //
  // Create an Exit Boot Services event.
  //
```

```
  Status = gBS->CreateEvent (
                 EFI_EVENT_SIGNAL_EXIT_BOOT_SERVICES,
                 EFI_TPL_NOTIFY,
                 AbcNotifyExitBootServices,
                 NULL,
                 &gExitBootServicesEvent
                 );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Create a Set Virtual Address Map event.
  //
  Status = gBS->CreateEvent (
                 EFI_EVENT_SIGNAL_SET_VIRTUAL_ADDRESS_MAP,
                 EFI_TPL_NOTIFY,
                 AbcNotifySetVirtualAddressMap,
                 NULL,
                 &gSetVirtualAddressMapEvent
                 );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Retrieve the Loaded Image Protocol from image handle
  //
  Status = gBS->OpenProtocol (
                 ImageHandle,
                 &gEfiLoadedImageProtocolGuid,
                 (VOID **)&LoadedImage,
                 ImageHandle,
                 ImageHandle,
                 EFI_OPEN_PROTOCOL_GET_PROTOCOL
                 );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Fill in the Unload() service of the Loaded Image Protocol
  //
  LoadedImage->Unload = AbcUnload;

  //
  // Perform additional driver initialization here
  //

  return EFI_SUCCESS;
}
```

Example 7-16. Runtime Driver Entry Point with Unload Feature

8
Private Context Data Structures

EFI drivers that manage more than one controller need to be designed with re-entrancy in mind. This means that the global variables should not be used to track information about individual controllers. Instead, data structures should be allocated with the EFI Memory Services for each controller, and those data structures should contain all the information that the driver requires to manage each individual controller. This chapter will introduce some object-oriented programming techniques that can be applied to drivers that manage controllers. These techniques can simplify the driver design and implementation. The concept of a private context data structure that contains all the information that is required to manage a controller will be introduced. This data structure contains the public data fields, public services, private data fields, and private services that an EFI driver may require to manage a controller.

Some classes of EFI drivers do not require the use of these data structures. If an EFI driver only produces a single protocol or it manages at most one device, then the techniques presented here are not required. An initializing driver does not produce any services and does not manage any devices, so it will not use this technique. A service driver that produces a single protocol and does not manage any devices likely will not use this technique. A root bridge driver that manages a single root bridge device likely will not use this technique, but a root bridge driver that manages more than one root bridge device should use this technique. Finally, all EFI drivers that follow the EFI Driver Model should use this technique. Even if the driver writer is convinced that the EFI driver will manage only a single device in a platform, this technique should still be used because it will simplify the process of updating the driver to manage more than one device. The driver writer should make as few device and platform assumptions as possible when designing a new driver.

It is possible to use other techniques to track the information that is required to manage multiple controllers in a re-entrant-safe manner, but those techniques will likely require more overhead in the driver itself to manage this information. The techniques presented here are intended to produce small driver executables, and these techniques are used throughout the *EFI 1.10 Sample Implementation*.

8.1 Containing Record Macro

The containing record macro, which is called **`CR()`**, enables good object-oriented programming practices. It returns a pointer to the structure using a pointer to one of the structure's fields. EFI drivers that produce protocols use this macro to retrieve the private context data structure from a pointer to a produced protocol interface. Protocol functions are required to pass in a pointer to the protocol instance as the first argument to the function. C++ does this automatically, and the pointer to the object instance is called a *this* pointer. Since EFI drivers are written in C, a close equivalent is implemented by requiring that the first argument of every protocol function be the pointer to the protocol's instance structure called *`This`*. Each protocol function then uses the **`CR()`** macro to retrieve a pointer to the private context data structure from this first argument called *`This`*.

Example 8-1 is the definition of the **CR()** macro. The **CR()** macro is provided a pointer to the following:

- A field in a data structure
- The name of the field

It uses this information to compute the offset of the field in the data structure and subtracts this offset from the pointer to the field. This calculation results in a pointer to the data structure that contains the specified field. **_CR()** returns a pointer to the data structure that contains the specified field. For debug builds, **CR()** also does an additional check to make sure the signature matches. If the signature does not match, then an **ASSERT()** message is generated and the system is haled. For production builds, the signature checks are skipped. Most EFI drivers define additional macros based on the **CR()** macro that retrieve the private context data structure based on a *This* pointer to a produced protocol.

```
#define _CR(Record, TYPE, Field)  \
  ((TYPE *) ( (CHAR8 *)(Record) - (CHAR8 *) &(((TYPE *) 0)->Field)))

#ifdef EFI_DEBUG

  #define CR(record, TYPE, field, signature)                          \
            _CR(record, TYPE, field)->Signature != signature ?        \
            (TYPE *) (_DEBUG_ASSERT("CR has Bad Signature"), record) : \
            _CR(record, TYPE, field)

#else

  #define CR(Record, TYPE, Field, Signature)   \
            _CR(Record, TYPE, Field)

#endif
```

Example 8-1. Containing Record Macro Definitions

8.2 Data Structure Design

Proper data structure design is one of the keys to making drivers both simple and easy to maintain. If a driver writer fails to include fields in a private context data structure, then it may require a complex algorithm to retrieve the required data through the various EFI services. By designing in the proper fields, these complex algorithms can be avoided and the driver will have a smaller executable footprint. Static data and commonly accessed data and services that are related to the management of a device should be placed in a private context data structure. Another requirement is that the private context data structure must be easy to find when an I/O service that is produced by the driver is called. The I/O services that are produced by a driver are exported through protocol interfaces, and all protocol interface include a *This* parameter as the first argument. The *This* parameter is a pointer to the protocol interface that contains the I/O service being called. The data structure design presented here will show how the *This* pointer that is passed into an I/O service can be used to very easily gain access to the private context data structure.

The following driver types will typically use private context data structures:

- Root bridge drivers
- Device drivers

- Bus drivers
- Hybrid drivers

Hybrid drivers may even use two different private context data structures—one for the bus controller and one for the child controllers it produces. A private context data structure is typically composed of the following types of fields:

- A signature for the data structure
- The handle of the controller or the child that is being managed or produced
- The group of protocol interfaces that are being consumed
- The group of protocol interfaces that are being produced
- Private data fields and services that are used to manage a specific controller

The signature is useful when debugging EFI drivers. These signatures are composed of four ASCII characters. When memory dumps are performed, these signatures stand out, so the beginning of specific data structures can be identified. Memory dump tools with search capabilities can also be used to find specific private context data structures in memory. In addition, debug builds of EFI drivers can perform signature checks whenever these private context data structures are accessed. If the signature does not match, then an **ASSERT()** can be generated. If one of these **ASSERT()** messages are observed, then an EFI driver was likely passed in a bad or corrupt memory pointer.

Device drivers will typically store the handle of the device they are managing in a private context data structure. This mechanism provides quick access to the device handle if it is needed during I/O operations or driver-related operations. Root bridge drivers and bus drivers will typically store the handle of the child that was created, and a hybrid driver will typically store both the handle of the bus controller and the handle of the child controller that was produced.

The group of consumed protocol interfaces is simply a set of pointers to the protocol interfaces that are opened in the **Start()** function of the driver's **EFI_DRIVER_BINDING_PROTOCOL**. As each protocol interface is opened using **gBS->OpenProtocol()**, a pointer to the consumed protocol interface is stored in the private context data structure. These same protocols must be closed in the **Stop()** function of the driver's **EFI_DRIVER_BINDING_PROTOCOL** with calls to **gBS->CloseProtocol()**.

The group of produced protocol interfaces declares the storage for the protocols that the driver produces. These protocols typically provide software abstractions for console or boot devices.

The number and types of the private data fields vary from driver to driver. These fields contain the context information for a device that is not contained in the consumed or produced protocols. For example, a driver for a disk device may store information about the geometry of the disk such as the number of cylinders, number of heads, and number of sectors on the physical disk that the driver is managing.

Example 8-2 below shows a generic template for a private context data structure that can be used for root bridge drivers, device drivers, bus drivers, or hybrid drivers. The **#define** statement at the beginning is used to initialize the *Signature* field when the private context data structure is allocated. This same **#define** statement is used to verify the *Signature* field whenever a driver accesses the private context data structure. The data structure itself contains the following:

- Signature
- Handle of the device being managed
- Pointers to the consumed protocols
- Storage for the produced protocols
- Any additional private data fields that are required to manage the device

The last part of this figure contains a set of macros that help retrieve a pointer to the private context data structure from a *This* pointer for each of the produced protocols. These macros are the simple mechanism that allows the private data fields to be accessed from the services in each of the produced protocols.

```
#define <<DRIVER_NAME>>_PRIVATE_DATA_SIGNATURE
EFI_SIGNATURE_32('A','B','C','D')

typedef struct {
  UINTN                   Signature;
  EFI_HANDLE              Handle;

  //
  // Pointers to consumed protocols
  //
  EFI_<<PROTOCOL_NAME_C1>>_PROTOCOL  *<<ProtocolNameC1>>;
  EFI_<<PROTOCOL_NAME_C2>>_PROTOCOL  *<<ProtocolNameC2>>;
  // . . .
  EFI_<<PROTOCOL_NAME_Cn>>_PROTOCOL  *<<ProtocolNameCn>>;

  //
  // Produced protocols
  //
  EFI_<<PROTOCOL_NAME_P1>>_PROTOCOL  <<ProtocolNameP1>>;
  EFI_<<PROTOCOL_NAME_P2>>_PROTOCOL  <<ProtocolNameP2>>;
  // . . .
  EFI_<<PROTOCOL_NAME_Pm>>_PROTOCOL  <<ProtocolNamePm>>;

  //
  // Private functions and data fields
  //
} <<DRIVER_NAME>>_PRIVATE_DATA;

#define <<DRIVER_NAME>_PRIVATE_DATA_FROM_<<PROTOCOL_NAME_P1>>_THIS(a) \
  CR(                                                                 \
    a,                                                                \
    <<DRIVER_NAME>>_PRIVATE_DATA,                                     \
    <<ProtocolNameP1>>,                                               \
    <<DRIVER_NAME>>_PRIVATE_DATA_SIGNATURE                            \
  )

#define <<DRIVER_NAME>_PRIVATE_DATA_FROM_<<PROTOCOL_NAME_P2>>_THIS(a) \
  CR(                                                                 \
    a,                                                                \
    <<DRIVER_NAME>>_PRIVATE_DATA,                                     \
    <<ProtocolNameP2>>,                                               \
    <<DRIVER_NAME>>_PRIVATE_DATA_SIGNATURE                            \
  )

// . . .

#define <<DRIVER_NAME>_PRIVATE_DATA_FROM_<<PROTOCOL_NAME_Pm>>_THIS(a) \
  CR(                                                                 \
    a,                                                                \
    <<DRIVER_NAME>>_PRIVATE_DATA,                                     \
    <<ProtocolNamePm>>,                                               \
    <<DRIVER_NAME>>_PRIVATE_DATA_SIGNATURE                            \
  )
```

Example 8-2. Private Context Data Structure Template

Example 8-3 below shows an example of the private context data structure from the disk I/O driver that is listed in Appendix D. It contains the **#define** statement for the data structure's signature. In this case, the signature is the ASCII string "dskI." It also contains a pointer to the only protocol that this driver consumes, which is the Block I/O Protocol, and it contains storage for the only protocol that this driver produces, which is the Disk I/O Protocol. It also has one private data field that caches the size of the blocks that the Block I/O Protocol supports. The macro at the bottom retrieves the private context data structure from a pointer to the *DiskIo* field.

```
#define DISK_IO_PRIVATE_DATA_SIGNATURE  EFI_SIGNATURE_32 ('d','s','k','I')

typedef struct {
  UINTN                   Signature;
  EFI_DISK_IO_PROTOCOL    DiskIo;
  EFI_BLOCK_IO_PROTOCOL   *BlockIo;
  UINT32                  BlockSize;
} DISK_IO_PRIVATE_DATA;

#define DISK_IO_PRIVATE_DATA_FROM_THIS(a) \
  CR (a, DISK_IO_PRIVATE_DATA, DiskIo, DISK_IO_PRIVATE_DATA_SIGNATURE)
```

Example 8-3. Simple Private Context Data Structure

Example 8-4 below shows a more complex private context data structure from the IDE bus driver, which is in the *EFI 1.10 Sample Implementation*. It contains the signature "ibid" and a *Handle* field that is the child handle for a disk device that the IDE bus driver produced. It also contains pointers to the consumed protocols, which are the Device Path Protocol and PCI I/O Protocol, and storage for the Block I/O Protocol that is produced by this driver. In addition, there are a large number of private data fields that are used during initialization, read operations, write operations, flush operations, and error recovery. Details on how these private fields are used can be found in the source code to the IDE bus driver in the *EFI 1.10 Sample Implementation*.

```
#define IDE_BLK_IO_DEV_SIGNATURE   EFI_SIGNATURE_32('i','b','i','d')

typedef struct {
    UINT32                        Signature;

    EFI_HANDLE                    Handle;
    EFI_BLOCK_IO_PROTOCOL         BlkIo;
    EFI_BLOCK_IO_MEDIA            BlkMedia;
    EFI_DEVICE_PATH_PROTOCOL      *DevicePath;
    EFI_PCI_IO_PROTOCOL           *PciIo;
    IDE_BUS_DRIVER_PRIVATE_DATA   *IdeBusDriverPrivateData;

    //
    // Local data for the IDE interface goes here
    //
    EFI_IDE_CHANNEL               Channel;
    EFI_IDE_DEVICE                Device;
    UINT16                        Lun;
    IDE_DEVICE_TYPE               Type;

    IDE_BASE_REGISTERS            *IoPort;
    UINT16                        AtapiError;

    INQUIRY_DATA                  *pInquiryData;
    IDENTIFY                      *pIdData;
    ATA_PIO_MODE                  PioMode;
```

```
    CHAR8                          ModelName[41];
    REQUEST_SENSE_DATA             *SenseData;
    UINT8                          SenseDataNumber;
    UINT8                          *Cache;

    EFI_UNICODE_STRING_TABLE       *ControllerNameTable;
} IDE_BLK_IO_DEV;

#define IDE_BLOCK_IO_DEV_FROM_THIS(a) \
  CR(a, IDE_BLK_IO_DEV, BlkIo, IDE_BLK_IO_DEV_SIGNATURE)
```

Example 8-4. Complex Private Context Data Structure

8.3 Allocating Private Context Data Structures

The private context data structures are allocated in the **Start()** function of the driver's **EFI_DRIVER_BINDING_PROTOCOL**. The service that is typically used to allocate the private context data structures is **gBS->AllocatePool()**. Example 8-5 below shows the generic template for allocating a private context data structure in the **Start()** function of the **EFI_DRIVER_BINDING_PROTOCOL**. This code example shows only a fragment from the **Start()** function. Chapter 9 covers the services that are produced by the **EFI_DRIVER_BINDING_PROTOCOL** in more detail.

```
EFI_STATUS
EFIAPI
<<DriverName>>DriverBindingStart (
  IN EFI_DRIVER_BINDING_PROTOCOL  *This,
  IN EFI_HANDLE                   ControllerHandle,
  IN EFI_DEVICE_PATH_PROTOCOL     *RemainingDevicePath  OPTIONAL
  )
{
  EFI_STATUS                     Status;
  <<DRIVER_NAME>>_PRIVATE_DATA   Private;

  //
  // Allocate the private context data structure
  //
  Status = gBS->AllocatePool (
                 EfiBootServicesData,
                 sizeof (<<DRIVER_NAME>>_PRIVATE_DATA),
                 (VOID**)&Private
                 );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Clear the contents of the allocated buffer
  //
  gBS->SetMem (Private, sizeof (<<DRIVER_NAME>>_PRIVATE_DATA), 0);
}
```

Example 8-5. Allocation of a Private Context Data Structure

Example 8-6 below shows the same generic template for the **Start()** function above except that it uses EFI Driver Library functions to allocate a private context data structure.

```
EFI_STATUS
EFIAPI
<<DriverName>>DriverBindingStart (
  IN EFI_DRIVER_BINDING_PROTOCOL  *This,
  IN EFI_HANDLE                   ControllerHandle,
  IN EFI_DEVICE_PATH_PROTOCOL     *RemainingDevicePath  OPTIONAL
  )
{
  EFI_STATUS                     Status;
  <<DRIVER_NAME>>_PRIVATE_DATA  Private;

  //
  // Allocate and clear the private context data structure
  //
  Private = EfiLibAllocateZeroPool (sizeof (<<DRIVER_NAME>>_PRIVATE_DATA));
  if (Private == NULL) {
    return EFI_OUT_OF_RESOURCES;
  }
}
```

Example 8-6. Library Allocation of Private Context Data Structure

Example 8-7 below shows a code fragment from the disk I/O driver that allocates the private context data structure for the disk I/O driver.

```
EFI_STATUS
EFIAPI
DiskIoDriverBindingStart (
  IN EFI_DRIVER_BINDING_PROTOCOL  *This,
  IN EFI_HANDLE                   ControllerHandle,
  IN EFI_DEVICE_PATH_PROTOCOL     *RemainingDevicePath  OPTIONAL
  )
{
  EFI_STATUS              Status;
  DISK_IO_PRIVATE_DATA    *Private;

  //
  // Initialize the disk I/O device instance.
  //
  Status = gBS->AllocatePool(
                 EfiBootServicesData,
                 sizeof (DISK_IO_PRIVATE_DATA),
                 &Private
                 );
  if (EFI_ERROR (Status)) {
    goto ErrorExit;
  }
  EfiZeroMem (Private, sizeof(DISK_IO_PRIVATE_DATA));
}
```

Example 8-7. Disk I/O Allocation of Private Context Data Structure

8.4 Freeing Private Context Data Structures

The private context data structures are freed in the **Stop()** function of the driver's **EFI_DRIVER_BINDING_PROTOCOL**. The service that is typically used to free the private context data structures is **gBS->FreePool()**. Example 8-8 below shows the generic template for allocating a private context data structure in the **Stop()** function of the **EFI_DRIVER_BINDING_PROTOCOL**. This code example shows only a fragment from the **Stop()** service. Chapter 9 covers the services that are produced by the **EFI_DRIVER_BINDING_PROTOCOL** in more detail.

```
EFI_STATUS
EFIAPI
<<DriverName>>DriverBindingStop (
  IN  EFI_DRIVER_BINDING_PROTOCOL  *This,
  IN  EFI_HANDLE                   ControllerHandle,
  IN  UINTN                        NumberOfChildren,
  IN  EFI_HANDLE                   *ChildHandleBuffer
  )
{
  EFI_STATUS                          Status;
  EFI_<<PROTOCOL_NAME_Pm>>_PROTOCOL   *<<ProtocolNamePm>>;
  <<DRIVER_NAME>>_PRIVATE_DATA        Private;

  //
  // Look up one of the driver's produced protocols
  //
  Status = gBS->OpenProtocol (
                 ControllerHandle,
                 &gEfi<<ProtocolNamePm>>ProtocolGuid,
                 &<<ProtocolNamePm>>,
                 This->DriverBindingHandle,
                 ControllerHandle,
                 EFI_OPEN_PROTOCOL_GET_PROTOCOL
                 );
  if (EFI_ERROR (Status)) {
    return EFI_UNSUPPORTED;
  }

  //
  // Retrieve the private context data structure from the produced protocol
  //
  Private = <<DRIVER_NAME>_PRIVATE_DATA_FROM_<<PROTOCOL_NAME_Pm>>_THIS (
              <<ProtocolNamePm>>
              );

  //
  // Free the private context data structure
  //
  Status = gBS->FreePool (Private);
  if (EFI_ERROR (Status)) {
    return Status;
  }
}
```

Example 8-8. Freeing a Private Context Data Structure

Example 8-9 below shows code fragments from the disk I/O driver that frees the private context data structure for the disk I/O driver.

```
EFI_STATUS
EFIAPI
DiskIoDriverBindingStop (
  IN  EFI_DRIVER_BINDING_PROTOCOL    *This,
  IN  EFI_HANDLE                     ControllerHandle,
  IN  UINTN                          NumberOfChildren,
  IN  EFI_HANDLE                     *ChildHandleBuffer
  )
{
  EFI_STATUS            Status;
  EFI_DISK_IO_PROTOCOL  *DiskIo;
  DISK_IO_PRIVATE_DATA  *Private;

  //
  // Get our context back.
  //
  Status = gBS->OpenProtocol (
                 ControllerHandle,
                 &gEfiDiskIoProtocolGuid,
                 &DiskIo,
                 This->DriverBindingHandle,
                 ControllerHandle,
                 EFI_OPEN_PROTOCOL_GET_PROTOCOL
                 );
  if (EFI_ERROR (Status)) {
    return EFI_UNSUPPORTED;
  }

  Private = DISK_IO_PRIVATE_DATA_FROM_THIS (DiskIo);

  gBS->FreePool (Private);
}
```

Example 8-9. Disk I/O Freeing of a Private Context Data Structure

8.5 Protocol Functions

The protocol functions that are produced by an EFI driver also need to access the private context data structure. These functions typically need access to the consumed protocols and the private data fields to perform the protocol function's required operation. Example 8-10 below shows a template for the implementation of a protocol function that retrieves the private context data structure using the **CR()** based macro and the *This* pointer for the produced protocol. The **Stop()** function from the **EFI_DRIVER_BINDING_PROTOCOL** uses the same **CR()** based macro to retrieve the private context data structure. The only difference is that the *This* pointer is not passed into the **Stop()** function. Instead, the **Stop()** function uses *ControllerHandle* to retrieve one of the produced protocols and then uses the **CR()** based macro with that protocol interface pointer to retrieve the private context data structure.

```
EFI_STATUS
EFIAPI
<<DriverName>><<ProtocolNamePn>><<FunctionNameM>> (
  IN EFI_<<PROTOCOL_NAME_PN>>_PROTOCOL *This,
  //
  // Additional function arguments here.
  //
  )
{
  <<DRIVER_NAME>>_PRIVATE_DATA  Private;

  //
  // Use This pointer to retrieve the private context structure
  //
  Private = <<DRIVER_NAME>_PRIVATE_DATA_FROM_<<PROTOCOL_NAME_Pn>>_THIS (This);
}
```

Example 8-10. Retrieving the Private Context Data Structure

Example 8-11 below shows a code fragment from the **ReadDisk()** service of the **EFI_DISK_IO_PROTOCOL** that is produced by the disk I/O driver. It uses the **CR()** macro and the *This* pointer to the **EFI_DISK_IO_PROTOCOL** to retrieve the **DISK_IO_PRIVATE_DATA** private context data structure.

```
EFI_STATUS
EFIAPI
DiskIoReadDisk (
  IN EFI_DISK_IO_PROTOCOL  *This,
  IN UINT32                MediaId,
  IN UINT64                Offset,
  IN UINTN                 BufferSize,
  OUT VOID                 *Buffer
  )
{
  EFI_STATUS            Status;
  DISK_IO_PRIVATE_DATA  *Private;

  Private = DISK_IO_PRIVATE_DATA_FROM_THIS (This);
}
```

Example 8-11. Retrieving the Disk I/O Private Context Data Structure

9
Driver Binding Protocol

The Driver Binding Protocol provides services that can be used to do the following:

- Connect a driver to a controller.
- Disconnect a driver from a controller.

EFI drivers that follow the EFI Driver Model are required to implement the Driver Binding Protocol. This requirement includes the following drivers:

- Device drivers
- Bus drivers
- Hybrid drivers

Root bridge driver, service drivers, and initializing drivers do not produce this protocol.

The Driver Binding Protocol is the most important protocol that a driver produces, because it is the one protocol that is used by the EFI Boot Services **gBS->ConnectController()** and **gBS->DisconnectController()**. These EFI Boot Services are used by the EFI boot manager to connect the console and boot devices that are required to boot an operating system. The implementation of the Driver Binding Protocol will vary depending upon the driver's class. Chapter 7 describes the various driver classes. The protocol interface structure for the Driver Binding Protocol is listed below for reference.

Protocol Interface Structure

```
typedef struct _EFI_DRIVER_BINDING_PROTOCOL {
  EFI_DRIVER_BINDING_SUPPORTED    Supported;
  EFI_DRIVER_BINDING_START        Start;
  EFI_DRIVER_BINDING_STOP         Stop;
  UINT32                          Version;
  EFI_HANDLE                      ImageHandle;
  EFI_HANDLE                      DriverBindingHandle;
} EFI_DRIVER_BINDING_PROTOCOL;
```

The Driver Binding Protocol contains the following three services:

- **Supported()**
- **Start()**
- **Stop()**

It also contains the following three data fields:

- *Version*
- *ImageHandle*
- *DriverBindingHandle*

The *ImageHandle* and *DriverBindingHandle* fields are typically preinitialized to **NULL**, and the EFI Driver Library functions automatically fill them in. The *Version* field does need to be initialized by the driver. Higher *Version* values signify a new driver. This field is a 32-bit value, but the values 0x0–0x0f and 0xfffffff0–0xffffffff are reserved for EFI drivers written by OEMs. IHVs may use the values 0x10–0xffffffef. Each time a new version of an EFI driver is released, the *Version* field should also be increased.

Many drivers use a *Version* value with a xx.xx.xx.xx revision scheme, so the driver can convey minor updates versus major updates. Drivers from third-party vendors should use this value to display the real true *Version* value that they use when referring to this driver. For example, if a driver has a version of 3.0.06, then the correct *Version* value would be 0x00030006. Whatever number scheme is chosen, it must be consistent with previously released drivers. Example 14-13 below shows how a Driver Binding Protocol is typically declared in a driver.

```
EFI_DRIVER_BINDING_PROTOCOL gAbcDriverBinding = {
  AbcDriverBindingSupported,    // Supported()
  AbcDriverBindingStart,        // Start()
  AbcDriverBindingStop,         // Stop()
  0x10,                         // Version
  NULL,                         // ImageHandle
  NULL                          // DriverBindingHandle
};
```

Example 9-1. Driver Binding Protocol Declaration

The **Supported()** service performs a quick check to see if a driver supports a controller. If the **Supported()** service passes, then the **Start()** service will be called to ask the driver to bind to a specific controller. The **Stop()** service does the opposite of the **Start()** service. It disconnects a driver from a controller and frees any resources that were allocated in the **Start()** services. None of these services are allowed to use any of the console I/O-related protocols. Instead, if an error condition is detected, then the error should be recorded and returned through the Driver Diagnostics Protocol. A driver may also use the **DEBUG()** and **ASSERT()** macros to send messages to the standard error console if it is active.

The implementations of the Driver Binding Protocol will change in complexity depending on the driver type. A device driver is fairly simple to implement. A bus driver or a hybrid driver may be more complex because it has to manage both the bus controller and the child controllers. These implementations will be discussed in the following sections.

The **EFI_DRIVER_BINDING_PROTOCOL** is installed onto the driver's image handle. It is possible for a driver to produce more than one instance of the Driver Binding Protocol. All additional instances of the Driver Binding Protocol should be installed onto new handles. The installation of the Driver Binding Protocol is handled by the EFI Library Service **EfiLibInstallAllDriverProtocols()**. This service is covered in more detail in section 7.1. If an error is generated in the installation of the Driver Binding Protocol, then the entire driver should fail.

The implementation of the Driver Binding Protocol for a specific driver is typically found in the file **<<DriverName>>.c**. This file contains the instance of the **EFI_DRIVER_BINDING_PROTOCOL** along with the implementation of the **Supported()**, **Start()**, and **Stop()** services. Example 9-2 below shows the template for the Driver Binding Protocol.

```
#define <<DRIVER_NAME>_VERSION 0x00000010

EFI_DRIVER_BINDING_PROTOCOL g<<DriverName>>DriverBinding = {
  <<DriverName>>DriverBindingSupported,  // Supported()
  <<DriverName>>DriverBindingStart,      // Start()
  <<DriverName>>DriverBindingStop,       // Stop()
  <<DRIVER_NAME>>_VERSION,               // Version
  NULL,                                  // ImageHandle
  NULL                                   // DriverBindingHandle
};

EFI_STATUS
<<DriverName>>DriverBindingSupported (
  IN EFI_DRIVER_BINDING_PROTOCOL  *This,
  IN EFI_HANDLE                   ControllerHandle,
  IN EFI_DEVICE_PATH_PROTOCOL     *RemainingDevicePath  OPTIONAL
  )
{
}

EFI_STATUS
<<DriverName>>DriverBindingStart (
  IN EFI_DRIVER_BINDING_PROTOCOL  *This,
  IN EFI_HANDLE                   ControllerHandle,
  IN EFI_DEVICE_PATH_PROTOCOL     *RemainingDevicePath  OPTIONAL
  )
{
}

EFI_STATUS
<<DriverName>>DriverBindingStop (
  IN EFI_DRIVER_BINDING_PROTOCOL  *This,
  IN EFI_HANDLE                   ControllerHandle,
  IN UINTN                        NumberOfChildren,
  IN EFI_HANDLE                   *ChildHandleBuffer
  )
{
}
```

Example 9-2. Driver Binding Protocol Template

The **Supported()**, **Start()**, and **Stop()** services are covered in detail in section 9.1 of the *EFI 1.10 Specification,* including the algorithms for implementing these services for device drivers and bus drivers. If a driver produces multiple instances of the Driver Binding Protocol, then they will be installed in the driver entry point. Each instance of the Driver Binding Protocol is implemented using the same guidelines. The different instances may share worker functions to reduce the size of the driver.

Once a Driver Binding Protocol is implemented, it can be tested with the following EFI Shell commands:

- `load`
- `connect`
- `disconnect`
- `reconnect`

The EFI boot manager will also test the services of the Driver Binding Protocol when the drivers for console devices and boot devices are connected so an operating system may be booted.

10
Component Name Protocol

The Component Name Protocol provides a human-readable name for drivers and the devices that drivers manage. This protocol applies only to EFI drivers that follow the EFI Driver Model, which includes the following:

- Device drivers
- Bus drivers
- Hybrid drivers

Root bridge drivers, service drivers, and initializing drivers do not produce this protocol. It is recommended that this protocol be implemented for all drivers because it provides an easy method for users to associate the various drivers to the controllers they manage. This protocol is also required by *DIG64*.

The Component Name Protocol provides human-readable names as null-terminated Unicode strings and can support one or more languages. At a minimum, the English language should be supported. The human-readable name that a driver produces is from that specific driver's perspective. If multiple drivers are managing the same controller handle, then there will be multiple instances of the Component Name Protocol that may provide different names for a controller. The consumers of the Component Name Protocol will have to decide which names to display when multiple names are available. For example, a PCI bus driver may produce a name for a PCI slot like "PCI Slot #2," and the driver for the SCSI adapter that is inserted into that same PCI slot may produce a name like "XYZ SCSI Host Controller." Both names describe the same physical device from each driver's perspective, and both names are useful depending on how they are used.

It is also suggested that these human-readable names be limited to about 40 Unicode characters in length so consumers of this protocol can easily fit these strings on standard console devices. The protocol interface structure for the Component Name Protocol is listed below for reference.

Protocol Interface Structure

```
typedef struct _EFI_COMPONENT_NAME_PROTOCOL {
  EFI_COMPONENT_NAME_GET_DRIVER_NAME       GetDriverName;
  EFI_COMPONENT_NAME_GET_CONTROLLER_NAME   GetControllerName;
  CHAR8                                    *SupportedLanguages;
} EFI_COMPONENT_NAME_PROTOCOL;
```

The Component Name Protocol advertises the languages it supports in a data field called `SupportedLanguages`. This data filed is a null-terminated ASCII string that contains one or more ISO 639-2 language codes. Each language code is composed of three ASCII characters. For example, English is specified by "eng," Spanish by "spa," and French by "fra." A consumer of the Component Name Protocol can parse the `SupportedLanguages` data field to see if the protocol supports a language in which the consumer is interested. This data field can also be used by the implementation of the Component Name Protocol to see if a requested language is supported.

The **GetDriverName()** service is used to retrieve the name of the driver, and the **GetControllerName()** service is used to retrieve the names of the controllers that the driver is currently managing or the names of the child controllers that the driver may have produced. The simplest implementation of the Component Name Protocol provides the name of the driver. The next most complex implementation is for a device driver that provides both the name of the driver and the names of the controllers that the driver is managing. The most complex implementation is for a bus driver or a hybrid driver that produces names for the driver, names for the bus controllers it is managing, and names for the child controllers that the driver has produced. All three of these implementations will be discussed in the sections that follow.

A few EFI Driver Library functions are provided to simplify the implementation of the Component Name Protocol. These library functions aid in registering and retrieving human-readable strings. Some drivers have fixed names for the drivers and controllers that they manage, and other drivers may create names on the fly based on information that is retrieved from the platform or the controller itself. The EFI Driver Library functions of interest are **EfiLibLookupUnicodeString()**, **EfiLibAddUnicodeString()**, and **EfiLibFreeUnicodeStringTable()**.

The **EFI_COMPONENT_NAME_PROTOCOL** must be installed onto the same handle as the **EFI_DRIVER_BINDING_PROTOCOL**. This install operation is done in the driver's entry point. The installation of the **EFI_DRIVER_BINDING_PROTOCOL** and the **EFI_COMPONENT_NAME_PROTOCOL** is covered in section 7.1. If an error is generated in the installation of the Component Name Protocol, then this error should not cause the entire driver to fail.

The implementation of the Component Name Protocol for a specific driver is typically found in the file **ComponentName.c**. This file contains the following:

- The instance of the **EFI_COMPONENT_NAME_PROTOCOL**
- A static table of driver names
- Static tables of controller names
- The implementation of the **GetDriverName()** and **GetControllerName()** services

For drivers that produce dynamic names for controller or children, the allocation and management of these dynamic names will be performed in the **Start()** and **Stop()** services of the **EFI_DRIVER_BINDING_PROTOCOL**. In addition, dynamic name tables require extra fields in the driver's private context data structure that point to the dynamic name tables. See chapter 8 for details on the design of private context data structures.

Once a Component Name Protocol is implemented, it can be tested with the EFI Shell commands **devices** and **drivers**. Platform vendors are also expected to implement extensions to the EFI boot manager that allow the user to manage the various devices in a platform. These EFI boot manager extensions should use the services of the Component Name Protocol to display the names of devices and the drivers that are managing those devices.

10.1 Driver Name

The simplest implementation of the **EFI_COMPONENT_NAME_PROTOCOL** produces a human-readable name for the driver itself and does not provide any names for the controllers or the children that the driver is managing. Example 10-1 below shows a template for this type of implementation. The instance of the **EFI_COMPONENT_NAME_PROTOCOL** is installed onto the driver handle in the driver's entry point. The **GetControllerName()** service always returns **EFI_UNSUPPORTED**, and the **GetDriverName()** service uses an EFI Driver Library function to retrieve the driver's name in the correct language from a static table of driver names. The static table of driver names contains two elements per entry. The first element is a three-character ASCII string that contains the ISO 639-2 language code for the driver name in the second element. The second element is a Unicode string that represents the name of the driver in the language specified by the first element. The static table is terminated by two **NULL** elements. Example 10-1 below shows an example with the three supported languages of English, Spanish, and French. Appendix D contains the source code to the disk I/O driver that produces the name of the disk I/O driver only in English.

```
//
// EFI Component Name Protocol
//
EFI_COMPONENT_NAME_PROTOCOL g<<DriverName>>ComponentName = {
  <<DriverName>>ComponentNameGetDriverName,
  <<DriverName>>ComponentNameGetControllerName,
  "engspafra”
};

//
// Static table of driver names in different languages
//
static EFI_UNICODE_STRING_TABLE m<<DriverName>>DriverNameTable[] = {
  { "eng", L"Insert English Driver Name Here" },
  { "spa", L"Insert Spanish Driver Name Here" },
  { "fra", L"Insert French Driver Name Here" },
  { NULL, NULL }
};

EFI_STATUS
<<DriverName>>ComponentNameGetDriverName (
  IN  EFI_COMPONENT_NAME_PROTOCOL  *This,
  IN  CHAR8                        *Language,
  OUT CHAR16                       **DriverName
  )
{
  return EfiLibLookupUnicodeString (
           Language,
           g<<DriverName>>ComponentName.SupportedLanguages,
           m<<DriverName>>DriverNameTable,
           DriverName
           );
}

EFI_STATUS
<<DriverName>>ComponentNameGetControllerName (
  IN  EFI_COMPONENT_NAME_PROTOCOL  *This,
  IN  EFI_HANDLE                   ControllerHandle,
  IN  EFI_HANDLE                   ChildHandle        OPTIONAL,
```

```
  IN  CHAR8                          *Language,
  OUT CHAR16                         **ControllerName
  )
{
  return EFI_UNSUPPORTED;
}
```

Example 10-1. Driver Name

10.2 Device Drivers

Device drivers that implement the Component Name Protocol will typically provide the name of the driver using the method described in section 10.1. This means that the implementation of the **GetDriverName()** service is always the same, and the only element that changes is the static table of driver names for each of the supported languages. The **GetControllerName()** service of the Component Name Protocol is where most of the work is performed.

Example 10-2 below shows the template for the **GetControllerName()** service for a device driver that produces static names for the controllers that it manages. The static controller names in different languages are declared in exactly the same way that the static driver names were declared in section 10.1. The **GetControllerName()** service needs to make more parameter checks than **GetDriverName()**, because it needs to make sure that the *ControllerHandle* and *ChildHandle* parameters match a controller that the driver is currently managing. Device drivers do not produce any child handles, so the *ChildHandle* parameter must be **NULL**. The *ControllerHandle* parameter must be checked to verify that it is a controller handle that the driver is currently managing. This check is done by opening a protocol on *ControllerHandle* that was opened **BY_DRIVER** in the driver's **Start()** service of the **EFI_DRIVER_BINDING_PROTOCOL**. If the status code returned is **EFI_ALREADY_STARTED**, then *ControllerHandle* is currently managed by the driver. If the status code returned is not **EFI_ALREADY_STARTED**, then *ControllerHandle* is not being managed by the driver.

Once a device driver has verified that *ChildHandle* is **NULL** and *ControllerHandle* represents a controller it is managing, then an EFI Driver Library Service can be used to retrieve the human-readable name of the controller from a static table of controller names in a specified language.

```
//
// Static table of Controller Names in different languages
//
static EFI_UNICODE_STRING_TABLE m<<DriverName>>ControllerNameTable[] = {
  { "eng", L"Insert English Controller Name Here " },
  { "spa", L"Insert Spanish Controller Name Here " },
  { "fra", L"Insert French Controller Name Here " },
  { NULL, NULL }
};

EFI_STATUS
<<DriverName>>ComponentNameGetControllerName (
  IN  EFI_COMPONENT_NAME_PROTOCOL  *This,
  IN  EFI_HANDLE                   ControllerHandle,
  IN  EFI_HANDLE                   ChildHandle        OPTIONAL,
  IN  CHAR8                        *Language,
  OUT CHAR16                       **ControllerName
```

```
  )
{
  EFI_STATUS                          Status;
  EFI_<<PROTOCOL_NAME_Cx>>_PROTOCOL  *<<ProtocolNameCx>>;

  //
  // ChildHandle must be NULL for device drivers
  //
  if (ChildHandle != NULL) {
    return EFI_UNSUPPORTED;
  }

  //
  // Make sure this driver is currently managing ControllerHandle
  //
  Status = gBS->OpenProtocol (
                 ControllerHandle,
                 &gEfi<<ProtocolNameCx>>ProtocolGuid,
                 (VOID **)&<<ProtocolNameCx>>,
                 g<<DriverName>>DriverBinding.DriverBindingHandle,
                 ControllerHandle,
                 EFI_OPEN_PROTOCOL_BY_DRIVER
                 );
  if (Status != EFI_ALREADY_STARTED) {
    gBS->CloseProtocol (
           ControllerHandle,
           &gEfi<<ProtocolNameCx>>ProtocolGuid,
           g<<DriverName>>DriverBinding.DriverBindingHandle,
           );
    return EFI_UNSUPPORTED;
  }

  //
  // Look up the controller name from a static table of controller names
  //
  return EfiLibLookupUnicodeString (
           Language,
           g<<DriverName>>ComponentName.SupportedLanguages,
           m<<DriverName>>ControllerNameTable,
           ControllerName
           );
}
```

Example 10-2. Device Driver with Static Controller Names

Some device drivers are able to extract information from the devices that they manage so they can provide more specific device names. The dynamic generation of controller names is more complex, but it can provide users with the detailed information they require to identify a specific device. For example, a driver for a disk device may be able to produce a static name such as “Hard Disk,” but a more specific name like “Seagate Barracuda ATA ST313620A Hard Disk” may be much more useful.

To support the dynamic generation of controller names, several additional steps must be taken. The first is that a pointer to the dynamic table of names must be added to the private context data structure for the controllers that a device driver manages. Example 10-3 below shows the addition of an **EFI_UNICODE_STRING_TABLE** field to the private context data structure discussed in chapter 8.

```
#define <<DRIVER_NAME>>_PRIVATE_DATA_SIGNATURE
EFI_SIGNATURE_32('A','B','C','D')

typedef struct {
  UINTN                             Signature;
  EFI_HANDLE                        Handle;

  //
  // Pointers to consumed protocols
  //
  EFI_<<PROTOCOL_NAME_C1>>_PROTOCOL  *<<ProtocolNameC1>>;
  EFI_<<PROTOCOL_NAME_C2>>_PROTOCOL  *<<ProtocolNameC2>>;
  // . . .
  EFI_<<PROTOCOL_NAME_Cn>>_PROTOCOL  *<<ProtocolNameCn>>;

  //
  // Produced protocols
  //
  EFI_<<PROTOCOL_NAME_P1>>_PROTOCOL  <<ProtocolNameP1>>;
  EFI_<<PROTOCOL_NAME_P2>>_PROTOCOL  <<ProtocolNameP2>>;
  // . . .
  EFI_<<PROTOCOL_NAME_Pm>>_PROTOCOL  <<ProtocolNamePm>>;

  //
  // Dynamically allocated table of controller names
  //
  EFI_UNICODE_STRING_TABLE           *ControllerNameTable;

  //
  // Additional Private functions and data fields
  //

}<<DRIVER_NAME>>_PRIVATE_DATA;
```

Example 10-3. Private Context Data Structure with a Dynamic Controller Name Table

The next update is to the **Start()** service of the **EFI_DRIVER_BINDING_PROTOCOL**. The **Start()** service needs to add a controller name in each supported language to *ControllerNameTable* in the private context data structure. The EFI Driver Library function **EfiLibAddUnicodeString()** can be used to add one or more names to a table. The *ContollerNameTable* must be initialized to **NULL** before the first name is added. Example 10-4 below shows an example of an English name being added to a dynamically allocated table of Unicode names. If more than one language is supported, then there would be an **EfiLibAddUnicodeString()** call for each language. The construction of the Unicode string for each language is not covered here. The format of names stored with devices varies depending on the bus type, and the translation from a bus-specific name format to a Unicode string cannot be standardized.

```
<<DRIVER_NAME>>_PRIVATE_DATA  *Private
CHAR16                        *ControllerName

//
// Get dynamic name from the device being managed
//

//
// Convert the device name to a Unicode string in a supported language
//

//
// Add the device name to the table of names stored in the private context data
// structure
//
EfiLibAddUnicodeString (
  "eng",
  g<<DriverName>>ComponentName.SupportedLanguages,
  &Private->ControllerNameTable,
  ControllerName
  );
```

Example 10-4. Adding a Controller Name to a Dynamic Controller Name Table

The **Stop()** service of the **EFI_DRIVER_BINDING_PROTOCOL** also needs to be updated. When a request is made for a driver to stop managing a controller, the table of controller names that were built in the **Start()** service must be freed. The EFI Driver Library function **EfiLibFreeUnicodeStringTable()** can be used to free the table of controller names. Example 10-5 below shows the code that should be added to the **Stop()** service. The private context data structure will be retrieved by the **Stop()** service so that the private context data structure can be freed. The call to **EfiLibFreeUnicodeStringTable()** should be made just before the private context data structure is freed.

```
<<DRIVER_NAME>>_PRIVATE_DATA  *Private

EfiLibFreeUnicodeStringTable (Private->ControllerNameTable);
```

Example 10-5. Freeing a Dynamic Controller Name Table

Finally, the **GetControllerName()** service becomes slightly more complex than the static controller name template. Because the table of controller names is now maintained in the private context data structure, the private context data structure needs to be retrieved based on the parameters passed into **GetControllerName()**. This retrieval is achieved by looking up a protocol that the driver has produced on *ControllerHandle* and using a pointer to that protocol and a **CR()** macro to retrieve a pointer to the private context data structure. Then, the private context data structure can be used with the EFI Driver Library function **EfiLibLookupUnicodeString()** to look up the controller's name in the dynamic table of controller names. Example 10-6 below shows a template for the **GetControllerName()** service that retrieves the controller name from a dynamic table that is stored in the private context data structure.

```
EFI_STATUS
<<DriverName>>ComponentNameGetControllerName (
  IN  EFI_COMPONENT_NAME_PROTOCOL  *This,
  IN  EFI_HANDLE                   ControllerHandle,
  IN  EFI_HANDLE                   ChildHandle        OPTIONAL,
  IN  CHAR8                        *Language,
  OUT CHAR16                       **ControllerName
  )
{
  EFI_STATUS                         Status;
  EFI_<<PROTOCOL_NAME_Cx>>_PROTOCOL  *<<ProtocolNameCx>>;
  EFI_<<PROTOCOL_NAME_Py>>_PROTOCOL  *<<ProtocolNamePy>>;
  <<DRIVER_NAME>>_PRIVATE_DATA       *Private;

  //
  // ChildHandle must be NULL for device drivers
  //
  if (ChildHandle != NULL) {
    return EFI_UNSUPPORTED;
  }

  //
  // Make sure this driver is currently managing ControllerHandle
  //
  Status = gBS->OpenProtocol (
                 ControllerHandle,
                 &gEfi<<ProtocolNameCx>>ProtocolGuid,
                 (VOID **)&<<ProtocolNameCx>>,
                 g<<DriverName>>DriverBinding.DriverBindingHandle,
                 ControllerHandle,
                 EFI_OPEN_PROTOCOL_BY_DRIVER
                 );
  if (Status != EFI_ALREADY_STARTED) {
    gBS->CloseProtocol (
           ControllerHandle,
           &gEfi<<ProtocolNameCx>>ProtocolGuid,
           g<<DriverName>>DriverBinding.DriverBindingHandle,
           );
    return EFI_UNSUPPORTED;
  }

  //
  // Retrieve an instance of a produced protocol from ControllerHandle
  //
  Status = gBS->OpenProtocol (
                 ControllerHandle,
                 &gEfi<<ProtocolNamePy>>ProtocolGuid,
                 (VOID **)&<<ProtocolNamePy>>,
                 g<<DriverName>>DriverBinding.DriverBindingHandle,
                 ControllerHandle,
                 EFI_OPEN_PROTOCOL_GET_PROTOCOL
                 );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Retrieve the private context data structure for ControllerHandle
  //
```

```
  Private = <<DRIVER_NAME>>_PRIVATE_DATA_FROM_<<PROTOCOL_NAME_Py>>_THIS (
              <<ProtocolNamePy>>
              );

  //
  // Look up the controller name from a dynamic table of controller names
  //
  return EfiLibLookupUnicodeString (
           Language,
           g<<DriverName>>ComponentName.SupportedLanguages,
           Private->ControllerNameTable,
           ControllerName
           );
}
```

Example 10-6. Device Driver with Dynamic Controller Names

10.3 Bus Drivers and Hybrid Drivers

There are many levels of support that a bus driver or hybrid driver may provide for the Component Name Protocol. These drivers can choose to provide a driver name as described in section 10.1. These drivers can also choose to provide names for the bus controllers that they manage and not provide any names for the children that they produce, such as the device drivers described in section 10.2. This section describes what bus drivers and hybrid drivers need to do to provide human-readable names for the child handles that they produce. The human-readable names for child handles can be provided through static or dynamic controller name tables.

Example 10-7 below shows an example of a driver that uses static name tables for both the bus controller and the child controllers. It first checks to make sure that *ControllerHandle* represents a bus controller that the driver is currently managing. Next it checks to see if *ChildHandle* is **NULL**. If *ChildHandle* is **NULL**, then the caller is asking for the human-readable name for the bus controller that the driver is managing. The **EfiLibLookupUnicodeString()** function is used to look up the human-readable name that is associated with *ControllerHandle*. If *ChildHandle* is not **NULL**, then the caller is requesting the name of a child controller that the driver has produced. First, a test is performed to make sure that *ChildHandle* is a handle that this driver produced. If that test succeeds, then the **EfiLibLookupUnicodeString()** function is used to look up the human-readable name that is associated with *ChildHandle*.

```
//
// Static table of controller names in different languages
//
static EFI_UNICODE_STRING_TABLE m<<DriverName>>ControllerNameTable[] = {
  { "eng", L"Insert English Bus Controller Name Here " },
  { "spa", L"Insert Spanish Bus Controller Name Here " },
  { "fra", L"Insert French Bus Controller Name Here " },
  { NULL, NULL }
};

//
// Static table of child names in different languages
//
static EFI_UNICODE_STRING_TABLE m<<DriverName>>ChildNameTable[] = {
```

```
  { "eng", L"Insert English Child Name Here " },
  { "spa", L"Insert Spanish Child Name Here " },
  { "fra", L"Insert French Child Name Here " },
  { NULL, NULL }
};

EFI_STATUS
<<DriverName>>ComponentNameGetControllerName (
  IN  EFI_COMPONENT_NAME_PROTOCOL  *This,
  IN  EFI_HANDLE                   ControllerHandle,
  IN  EFI_HANDLE                   ChildHandle        OPTIONAL,
  IN  CHAR8                        *Language,
  OUT CHAR16                       **ControllerName
  )
{
  EFI_STATUS                           Status;
  EFI_<<PROTOCOL_NAME_Cx>>_PROTOCOL    *<<ProtocolNameCx>>;
  EFI_OPEN_PROTOCOL_INFORMATION_ENTRY  *OpenInfoBuffer;
  UINTN                                EntryCount;
  UINTN                                Index;

  //
  // Make sure this driver is currently managing ControllerHandle
  //
  Status = gBS->OpenProtocol (
                 ControllerHandle,
                 &gEfi<<ProtocolNameCx>>ProtocolGuid,
                 (VOID **)&<<ProtocolNameCx>>,
                 g<<DriverName>>DriverBinding.DriverBindingHandle,
                 ControllerHandle,
                 EFI_OPEN_PROTOCOL_BY_DRIVER
                 );
  if (Status != EFI_ALREADY_STARTED) {
    gBS->CloseProtocol (
           ControllerHandle,
           &gEfi<<ProtocolNameCx>>ProtocolGuid,
           g<<DriverName>>DriverBinding.DriverBindingHandle,
           );
    return EFI_UNSUPPORTED;
  }

  if (ChildHandle == NULL) {
    //
    // Look up the controller name from a static table of controller names
    //
    return EfiLibLookupUnicodeString (
             Language,
             g<<DriverName>>ComponentName.SupportedLanguages,
             m<<DriverName>>ControllerNameTable,
             ControllerName
             );
  }

  //
  // Retrieve the list of agents that are consuming one of the protocols
  // on ControllerHandle that the children consume
  //
  Status = gBS->OpenProtocolInformation (
                 ControllerHandle,
```

```
                    &gEfi<<ProtocolNameCx>>ProtocolGuid,
                    &OpenInfoBuffer,
                    &EntryCount
                    );
  if (EFI_ERROR (Status)) {
    return EFI_UNSUPPORTED;
  }

  //
  // See if one of the agents is ChildHandle
  //
  Status = EFI_UNSUPPORTED;
  for (Index = 0; Index < EntryCount; Index++) {
    if (OpenInfoBuffer[Index].ControllerHandle == ChildHandle &&
        OpenInfoBuffer[Index].Attributes&EFI_OPEN_PROTOCOL_BY_CHILD_CONTROLLER)
{
      Status = EFI_SUCCESS;
    }
  }

  //
  // Free the information buffer
  //
  gBS->FreePool (OpenInfoBuffer);

  //
  // If ChildHandle was not one of the agents, then return EFI_UNSUPPORTED
  //
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Look up the child name from a static table of child names
  //
  return EfiLibLookupUnicodeString (
           Language,
           g<<DriverName>>ComponentName.SupportedLanguages,
           m<<DriverName>>ChildNameTable,
           ControllerName
           );
}
```

Example 10-7. Bus Driver with Static Controller and Child Names

The static tables for the controller names and the child names can be substituted with dynamic tables. This substitution requires the private context structure to be updated along with the **`Start()`** and **`Stop()`** services of the **`EFI_DRIVER_BINDING_PROTOCOL`**. Section 10.2 covers how this update is done for the controller names. Example 10-8 below shows how this update is performed for a bus driver that does not produce names for the bus controllers that it manages but does produce names for the child handles it creates.

```
EFI_STATUS
<<DriverName>>ComponentNameGetControllerName (
  IN  EFI_COMPONENT_NAME_PROTOCOL  *This,
  IN  EFI_HANDLE                   ControllerHandle,
  IN  EFI_HANDLE                   ChildHandle        OPTIONAL,
  IN  CHAR8                        *Language,
```

```
  OUT CHAR16                        **ControllerName
  )
{
  EFI_STATUS                            Status;
  EFI_<<PROTOCOL_NAME_Cx>>_PROTOCOL     *<<ProtocolNameCx>>;
  EFI_OPEN_PROTOCOL_INFORMATION_ENTRY   *OpenInfoBuffer;
  UINTN                                 EntryCount;
  UINTN                                 Index;
  EFI_<<PROTOCOL_NAME_Py>>_PROTOCOL     *<<ProtocolNamePy>>;
  <<DRIVER_NAME>>_PRIVATE_DATA          *Private;

  //
  // This version cannot return controller names, so ChildHandle
  // must not be NULL
  //
  if (ChildHandle == NULL) {
    return EFI_UNSUPPORTED;
  }

  //
  // Make sure this driver is currently managing ControllerHandle
  //
  Status = gBS->OpenProtocol (
                 ControllerHandle,
                 &gEfi<<ProtocolNameCx>>ProtocolGuid,
                 (VOID **)&<<ProtocolNameCx>>,
                 g<<DriverName>>DriverBinding.DriverBindingHandle,
                 ControllerHandle,
                 EFI_OPEN_PROTOCOL_BY_DRIVER
                 );
  if (Status != EFI_ALREADY_STARTED) {
    gBS->CloseProtocol (
           ControllerHandle,
           &gEfi<<ProtocolNameCx>>ProtocolGuid,
           g<<DriverName>>DriverBinding.DriverBindingHandle,
           );
    return EFI_UNSUPPORTED;
  }

  //
  // Retrieve the list of agents that are consuming one of the protocols
  // on ControllerHandle that the children consume
  //
  Status = gBS->OpenProtocolInformation (
                 ControllerHandle,
                 &gEfi<<ProtocolNameCx>>ProtocolGuid,
                 &OpenInfoBuffer,
                 &EntryCount
                 );
  if (EFI_ERROR (Status)) {
    return EFI_UNSUPPORTED;
  }

  //
  // See if one of the agents is ChildHandle
  //
  Status = EFI_UNSUPPORTED;
  for (Index = 0; Index < EntryCount; Index++) {
    if (OpenInfoBuffer[Index].ControllerHandle == ChildHandle &&
```

```
        OpenInfoBuffer[Index].Attributes&EFI_OPEN_PROTOCOL_BY_CHILD_CONTROLLER)
{
      Status = EFI_SUCCESS;
    }
  }

  //
  // Free the information buffer
  //
  gBS->FreePool (OpenInfoBuffer);

  //
  // If ChildHandle was not one of the agents, then return EFI_UNSUPPORTED
  //
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Retrieve the instance of a produced protocol from ChildHandle
  //
  Status = gBS->OpenProtocol (
                  ChildHandle,
                  &gEfi<<ProtocolNamePy>>ProtocolGuid,
                  (VOID **)&<<ProtocolNamePy>>,
                  g<<DriverName>>DriverBinding.DriverBindingHandle,
                  ChildHandle,
                  EFI_OPEN_PROTOCOL_GET_PROTOCOL
                  );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Retrieve the private context data structure for ChildHandle
  //
  Private = <<DRIVER_NAME>>_PRIVATE_DATA_FROM_<<PROTOCOL_NAME_Py>>_THIS (
              <<ProtocolNamePy>>
              );

  //
  // Look up the controller name from a dynamic table of child controller names
  //
  return EfiLibLookupUnicodeString (
           Language,
           g<<DriverName>>ComponentName.SupportedLanguages,
           Private->ControllerNameTable,
           ControllerName
           );
}
```

Example 10-8. Bus Driver with Dynamic Child Names

11
Driver Configuration Protocol

The Driver Configuration Protocol allows users to set the configuration options for the devices that drivers manage. This protocol applies only to EFI drivers that follow the EFI Driver Model, which includes the following:

- Device drivers
- Bus drivers
- Hybrid drivers

Root bridge drivers, service drivers, and initializing drivers do not produce this protocol. If a driver requires configuration information, the Driver Configuration Protocol must be implemented. *DIG64* requires that this protocol be implemented in the EFI drivers that own hardware devices.

The Driver Configuration Protocol provides a user interface for the user to configure devices in one or more languages. At a minimum, the English language should be supported. If multiple drivers are managing the same controller handle, then there may be multiple Driver Configuration Protocols present for that controller handle. The consumers of the Driver Configuration Protocol will have to decide how the multiple drivers that support configuration are presented to the user. For example, a PCI bus driver may produce a mechanism to enable or disable a specific slot, and the driver for the SCSI adapter that is inserted into that same PCI slot may produce configuration options for the SCSI host controller. Both set of configuration options may be useful to the user when managing the platform. The protocol interface structure for the Driver Configuration Protocol is listed below for reference.

Protocol Interface Structure

```
typedef struct _EFI_DRIVER_CONFIGURATION_PROTOCOL {
  EFI_DRIVER_CONFIGURATION_SET_OPTIONS     SetOptions;
  EFI_DRIVER_CONFIGURATION_OPTIONS_VALID   OptionsValid;
  EFI_DRIVER_CONFIGURATION_FORCE_DEFAULTS  ForceDefaults;
  CHAR8                                    *SupportedLanguages;
} EFI_DRIVER_CONFIGURATION_PROTOCOL;
```

The Driver Configuration Protocol advertises the languages it supports in a data field called `SupportedLanguages`. This data filed is a null-terminated ASCII string that contains one or more ISO 639-2 language codes. Each language code is composed of 3 ASCII characters. For example, English is specified by “eng,” Spanish by “spa,” and French by “fra.” A consumer of the Driver Configuration Protocol can parse the `SupportedLanguages` data field to see if the protocol supports a language in which the consumer is interested. The implementation of the Driver Configuration Protocol can also use this data field to see if a requested language is supported.

The **SetOptions()** service uses the console-related protocols such as the Simple Input Protocol and the Simple Text Output Protocol to allow the user to adjust the configuration options for a specific device that is being managed by the driver. This service is the only way that the driver is allowed to interact with the user through the console device. The driver may provide the typical banner type information here and provide an interactive menu if required. The design of the user interface should target the worst-case console configuration, which is a serial terminal such as VT-100 at 9600 baud. When choosing input keys for the user interface, the worst-case console configuration should be considered. See section 5.16 for suggestions on selecting input keys.

The **OptionsValid()** service is used to check if the current set of configuration options are valid for a specific device, and **ForceDefaults()** is an optional service that provides a mechanism to place a device in a default configuration.

The implementations of the Driver Configuration Protocol will vary in complexity depending on the driver type. A device driver is fairly simple to implement. A bus driver or a hybrid driver may be more complex because it may have to provide configuration settings for both the bus controller and the child controllers. Both of these implementations will be discussed in the following sections.

The **EFI_DRIVER_CONFIGURATION_PROTOCOL** must be installed onto the same handle as the **EFI_DRIVER_BINDING_PROTOCOL**. This install operation is done in the driver's entry point. Installing the **EFI_DRIVER_BINDING_PROTOCOL** and the **EFI_DRIVER_CONFIGURATION_PROTOCOL** is covered in section 7.1. If an error is generated when installing the Driver Configuration Protocol, then this error should not cause the entire driver to fail.

The implementation of the Driver Configuration Protocol for a specific driver is typically found in the file **DriverConfiguration.c**. This file contains the instance of the **EFI_DRIVER_CONFIGURATION_PROTOCOL** along with the implementation of the **SetOptions()**, **OptionsValid()**, and **ForceDefault()** services. Example 11-1 below shows a template for the implementation of the Driver Configuration Protocol.

```
EFI_DRIVER_CONFIGURATION_PROTOCOL g<<DriverName>>DriverConfiguration = {
  <<DriverName>>DriverConfigurationSetOptions,
  <<DriverName>>DriverConfigurationOptionsValid,
  <<DriverName>>DriverConfigurationForceDefaults,
  "eng"
};

EFI_STATUS
<<DriverName>>DriverConfigurationSetOptions (
  IN  EFI_DRIVER_CONFIGURATION_PROTOCOL         *This,
  IN  EFI_HANDLE                                ControllerHandle,
  IN  EFI_HANDLE                                ChildHandle  OPTIONAL,
  IN  CHAR8                                     *Language,
  OUT EFI_DRIVER_CONFIGURATION_ACTION_REQUIRED  *ActionRequired
  )
{
}

EFI_STATUS
<<DriverName>>DriverConfigurationOptionsValid (
  IN  EFI_DRIVER_CONFIGURATION_PROTOCOL  *This,
  IN  EFI_HANDLE                         ControllerHandle,
  IN  EFI_HANDLE                         ChildHandle  OPTIONAL
```

```
  )
{
}

EFI_STATUS
<<DriverName>>DriverConfigurationForceDefaults (
  IN  EFI_DRIVER_CONFIGURATION_PROTOCOL         *This,
  IN  EFI_HANDLE                                ControllerHandle,
  IN  EFI_HANDLE                                ChildHandle  OPTIONAL,
  IN  UINT32                                    DefaultType,
  OUT EFI_DRIVER_CONFIGURATION_ACTION_REQUIRED  *ActionRequired
  )
{
}
```

Example 11-1. Driver Configuration Protocol Template

Once a Driver Configuration Protocol is implemented, it can be tested with the EFI Shell command **drvcfg**. Platform vendors are also expected to implement extensions to the EFI boot manager that allow the user to configure the various devices in a platform. These EFI boot manager extensions use the services of the Component Name Protocol to display the names of devices and it will use the services of the Driver Configuration Protocol to allow the user to adjust the configuration settings for each device. The platform vendor also has the ability to validate the configuration of all the devices in the system prior to booting. In addition, the devices can be reset to their default configurations. Forcing default configurations may be automated if the platform firmware detects a corrupt configuration, or the platform vendor may allow the user to select a menu item to force defaults on a specific device or all devices at once.

The Driver Configuration Protocol is available only for devices that a driver is currently managing. Because EFI supports connecting the minimum number of drivers and devices that are required to establish consoles and gain access to the boot device, there may be many configurable devices that cannot be configured. As a result, when the user wishes to enter a “platform configuration” mode, the EFI boot manager will be required to connect all drivers to all devices, so the user will be able to see all the configurable devices in the platform.

11.1 Device Drivers

Device drivers will use the template from Example 11-1. Example 11-2 below shows the additional code that should be added to be beginning of each of the services of the Driver Configuration Protocol. The first step is to evaluate *ChildHandle*. If it is not **NULL**, then a request is being made to configure a child device. Because device drivers do not produce children, *ChildHandle* must be **NULL**. The next check is to make sure that *ControllerHandle* represents a device that the device driver is currently managing. This check is done by attempting to open a protocol that was opened **BY_DRIVER** in the **Start()** service. If the return code is not **EFI_ALREADY_STARTED**, then *ControllerHandle* is not being managed by this driver, so **EFI_UNSUPPORTED** is returned. The final step is to retrieve a protocol that was produced by the device driver from *ControllerHandle* and then use a **CR()** macro to retrieve the private context data structure for the device being managed.

```
  EFI_STATUS                        Status;
  EFI_<<PROTOCOL_NAME_Cx>>_PROTOCOL  *<<ProtocolNameCx>>;
```

```
  EFI_<<PROTOCOL_NAME_Py>>_PROTOCOL  *<<ProtocolNamePy>>;
  <<DRIVER_NAME>>_PRIVATE_DATA       *Private;

  //
  // Child handle must be NULL for a device driver
  //
  if (ChildHandle != NULL) {
    return EFI_UNSUPPORTED;
  }

  //
  // Make sure this driver is currently managing ControllerHandle
  //
  Status = gBS->OpenProtocol (
                 ControllerHandle,
                 &gEfi<<ProtocolNameCx>>ProtocolGuid,
                 (VOID **)&<<ProtocolNameCx>>,
                 g<<DriverName>>DriverBinding.DriverBindingHandle,
                 ControllerHandle,
                 EFI_OPEN_PROTOCOL_BY_DRIVER
                 );
  if (Status != EFI_ALREADY_STARTED) {
    gBS->CloseProtocol (
           ControllerHandle,
           &gEfi<<ProtocolNameCx>>ProtocolGuid,
           g<<DriverName>>DriverBinding.DriverBindingHandle,
           );
    return EFI_UNSUPPORTED;
  }

  //
  // Retrieve the an instance of a produced protocol from ControllerHandle
  //
  Status = gBS->OpenProtocol (
                 ControllerHandle,
                 &gEfi<<ProtocolNamePy>>ProtocolGuid,
                 (VOID **)&<<ProtocolNamePy>>,
                 g<<DriverName>>DriverBinding.DriverBindingHandle,
                 ControllerHandle,
                 EFI_OPEN_PROTOCOL_GET_PROTOCOL
                 );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Retrieve the private context data structure for ControllerHandle
  //
  Private = <<DRIVER_NAME>>_PRIVATE_DATA_FROM_<<PROTOCOL_NAME_Py>>_THIS (
              <<ProtocolNamePy>>
              );
```

Example 11-2. Retrieving the Private Context Data Structure

The implementation of **`SetOptions()`** will use the Console I/O Services provided by the Simple Input Protocol and the Simple Text Output Protocol to interact with the user. These protocols are described in sections 10.2 and 10.3 of the *EFI 1.10 Specification*. More advanced user interfaces may choose to use the services of the UGA Draw Protocol and the Simple Pointer Protocols

described in sections 10.5 and 10.10 of the *EFI 1.10 Specification*. **SetOptions()** shall not directly access a serial port, keyboard controller, or a VGA controller. The Simple Input Protocol is used to access the console input device that is available from the EFI System Table through **gST->ConIn**. The Simple Text Output Protocol is used to access the console output device that is available from the EFI System Table through **gST->ConOut**. The specific design of the user interface and the methods that are used to support multiple languages are not covered here. These design decisions are up to the individual driver writer.

SetOptions() stores its configuration information in nonvolatile storage. This configuration information should be stored with the device, so that the configuration information travels with the device if it is moved between platforms. The exact method for retrieving and storing configuration information on a device is device specific and will not be covered here. Typically, drivers will use the services of a bus I/O protocol to access the resources on a device to retrieve and store configuration information. For example, if a PCI controller has a flash device attached to it, the management of that flash device may be exposed through I/O or memory-mapped I/O registers described in the BARs associated with the PCI device. A PCI device driver can use the **Io.Read()**, **Io.Write()**, **Mem.Read()**, or **Mem.Write()** services of the PCI I/O Protocol to access the flash contents to retrieve and store configuration settings. Devices that are integrated onto the motherboard or are part of a FRU may use the EFI Variable Services such as **gRT->GetVariable()** and **gRT->SetVariable()** to store configuration information.

The **OptionsValid()** service will not interact with the user. Instead, it will read the device's current configuration information and make sure that the information contains a valid set of configuration options. If the configuration information cannot be retrieved or if the configuration information appears to be corrupt, then an error is returned.

The **ForceDefault()** service is optional and may simply return **EFI_UNSUPPORTED**. If this service is implemented, then it stores a valid set of default configuration settings in nonvolatile storage.

The implementation of the **SetOptions()** service could use **OptionsValid()** and **ForceDefaults()** to make sure the current options are valid and, if they are not valid, place the controller in its default configuration. The code fragment in Example 11-3 below shows the code that could be added to **SetOptions()** to implement this feature.

```
EFI_STATUS                                  Status;
EFI_DRIVER_CONFIGURATION_ACTION_REQUIRED  ActionRequired;

Status = This->OptionsValid (This, ControllerHandle, ChildHandle);
if (EFI_ERROR (Status)) {
  Status = This->ForceDefaults (
                   This,
                   ControllerHandle,
                   ChildHandle,
                   EFI_DRIVER_CONFIGURATION_SAFE_DEFAULTS,
                   &ActionRequired
                   );
}
```

Example 11-3. Validating Options in SetOptions()

11.2 Bus Drivers and Hybrid Drivers

A bus driver or hybrid driver may provide many levels of support for the Driver Configuration Protocol. These drivers can manage configuration settings for the bus controllers that they manage and not provide any configuration settings for the children that they produce. This choice means that they behave exactly like the device drivers described in section 11.1. This section describes what bus drivers and hybrid drivers need to do to manage configuration settings for the child handles that they produce.

Example 11-4 below shows the additional code that should be added to the beginning of each of the services of the Driver Configuration Protocol. The first step is to evaluate *ChildHandle*. If it is **NULL**, then a request is being made to configure the bus controller. If it is not **NULL**, then a request is being made to configure a child. The *ControllerHandle* must always be checked to make sure it represents a device that the driver is currently managing. This check is done by attempting to open a protocol that was opened **BY_DRIVER** in the **Start()** service. If the return code is not **EFI_ALREADY_STARTED**, then *ControllerHandle* is not being managed by this driver, so **EFI_UNSUPPORTED** is returned. If *ChildHandle* is not **NULL**, then an additional check must be made to verify that *ChildHandle* was produced by this driver. If that check passes, then a protocol that was produced on the *ChildHandle* is retrieved and the **CR()** macro is used to retrieve the private context data structure for *ChildHandle*.

```
EFI_STATUS                          Status;
EFI_<<PROTOCOL_NAME_Cx>>_PROTOCOL   *<<ProtocolNameCx>>;
EFI_<<PROTOCOL_NAME_Py>>_PROTOCOL   *<<ProtocolNamePy>>;
EFI_<<PROTOCOL_NAME_Pz>>_PROTOCOL   *<<ProtocolNamePz>>;
EFI_OPEN_PROTOCOL_INFORMATION_ENTRY *OpenInfoBuffer;
UINTN                               EntryCount;
UINTN                               Index;
<<DRIVER_NAME>>_PRIVATE_DATA        *Private;

//
// Make sure this driver is currently managing ControllerHandle
//
Status = gBS->OpenProtocol (
                ControllerHandle,
                &gEfi<<ProtocolNameCx>>ProtocolGuid,
                (VOID **)&<<ProtocolNameCx>>,
                g<<DriverName>>DriverBinding.DriverBindingHandle,
                ControllerHandle,
                EFI_OPEN_PROTOCOL_BY_DRIVER
                );
if (Status != EFI_ALREADY_STARTED) {
  gBS->CloseProtocol (
         ControllerHandle,
         &gEfi<<ProtocolNameCx>>ProtocolGuid,
         g<<DriverName>>DriverBinding.DriverBindingHandle,
         );
  return EFI_UNSUPPORTED;
}

//
// If child handle is NULL, then the bus controller is being configured
//
if (ChildHandle == NULL) {
  //
```

```
    // Retrieve the an instance of a produced protocol from ControllerHandle
    //
    Status = gBS->OpenProtocol (
                    ControllerHandle,
                    &gEfi<<ProtocolNamePy>>ProtocolGuid,
                    (VOID **)&<<ProtocolNamePy>>,
                    g<<DriverName>>DriverBinding.DriverBindingHandle,
                    ControllerHandle,
                    EFI_OPEN_PROTOCOL_GET_PROTOCOL
                   );
    if (EFI_ERROR (Status)) {
      return Status;
    }

    //
    // Retrieve the private context data structure for ControllerHandle
    //
    Private = <<DRIVER_NAME>>_PRIVATE_DATA_FROM_<<PROTOCOL_NAME_Py>>_THIS (
                <<ProtocolNamePy>>
                );

    //
    // Add user interface code to configure the bus controller here
    //

    return EFI_SUCCESS;
  }

  //
  // Retrieve the list of agents that are consuming one of the protocols
  // on ControllerHandle that the children opened BY_CHILD_CONTROLLER
  //
  Status = gBS->OpenProtocolInformation (
                  ControllerHandle,
                  &gEfi<<ProtocolNameCx>>ProtocolGuid,
                  &OpenInfoBuffer,
                  &EntryCount
                  );
  if (EFI_ERROR (Status)) {
    return EFI_UNSUPPORTED;
  }

  //
  // See if one of the agents is ChildHandle
  //
  Status = EFI_UNSUPPORTED;
  for (Index = 0; Index < EntryCount; Index++) {
    if (OpenInfoBuffer[Index].ControllerHandle == ChildHandle &&
        OpenInfoBuffer[Index].Attributes&EFI_OPEN_PROTOCOL_BY_CHILD_CONTROLLER)
{
      Status = EFI_SUCCESS;
    }
  }

  //
  // Free the information buffer
  //
  gBS->FreePool (OpenInfoBuffer);
```

```
  //
  // If ChildHandle was not one of the agents, then return EFI_UNSUPPORTED
  //
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Retrieve the instance of a produced protocol from ControllerHandle
  //
  Status = gBS->OpenProtocol (
                  ChildHandle,
                  &gEfi<<ProtocolNamePz>>ProtocolGuid,
                  (VOID **)&<<ProtocolNamePz>>,
                  g<<DriverName>>DriverBinding.DriverBindingHandle,
                  ChildHandle,
                  EFI_OPEN_PROTOCOL_GET_PROTOCOL
                 );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Retrieve the private context data structure for ControllerHandle
  //
  Private = <<DRIVER_NAME>>_PRIVATE_DATA_FROM_<<PROTOCOL_NAME_Pz>>_THIS (
              <<ProtocolNamePz>>
              );

  //
  // Add user interface code to configure the child controller here
  //

  return EFI_SUCCESS;
}
```

Example 11-4. Retrieving the Private Context Data Structure

11.3 Implementing SetOptions() as an Application

One possible design of the Driver Configuration Protocol is to implement the user interface as an EFI application that is stored with the device or in an EFI system partition. The implementation of the **`SetOptions()`** service would be changed so it does not produce a user interface. It would perform the same parameter checks as before and it would likely still retrieve the private context data structure. Then, instead of producing a user interface, it would use the **`gBS->LoadImage()`** and **`gBS->StartImage()`** services to load and execute the EFI application that provides the user interface. This application can then either directly update the new configuration settings, or it can pass the new configuration settings back to **`SetOptions()`**.

12
Driver Diagnostics Protocol

The Driver Diagnostics Protocol allows diagnostics to be executed on the devices that drivers manage. This protocol applies only to EFI drivers that follow the EFI Driver Model, which includes the following:

- Device drivers
- Bus drivers
- Hybrid drivers

Root bridge drivers, service drivers, and initializing drivers do not produce this protocol. If a driver implements diagnostic services, then the Driver Diagnostics Protocol must be implemented. *DIG64* requires that this protocol be implemented in the EFI drivers that own hardware devices.

The Driver Diagnostics Protocol provides diagnostics results in one or more languages. At a minimum, the English language should be supported. If multiple drivers are managing the same controller handle, then there may be multiple instances of Driver Diagnostics Protocol present for that controller handle. The consumers of the Driver Diagnostics Protocol will have to decide how the multiple drivers that support diagnostics are presented to the user. For example, a PCI bus driver may produce a mechanism to verify the functionality of a specific PCI slot, and the driver for a SCSI adapter that is inserted into that same PCI slot may produce diagnostics for the SCSI host controller. Both sets of diagnostics may be useful to the user when testing the platform. The protocol interface structure for the Driver Diagnostics Protocol is listed below for reference.

Protocol Interface Structure

```
typedef struct _EFI_DRIVER_DIAGNOSTICS_PROTOCOL {
  EFI_DRIVER_DIAGNOSTICS_RUN_DIAGNOSTICS  RunDiagnostics;
  CHAR8                                   *SupportedLanguages;
} EFI_DRIVER_DIAGNOSTICS_PROTOCOL;
```

The Driver Diagnostics Protocol advertises the languages it supports in a data field called `SupportedLanguages`. This data filed is a null-terminated ASCII string that contains one or more ISO 639-2 language codes. Each language code is composed of three ASCII characters. For example, English is specified by “eng,” Spanish by “spa,” and French by “fra.” A consumer of the Driver Diagnostics Protocol can parse the `SupportedLanguages` data field to see if the protocol supports a language in which the consumer is interested. This data field can also be used by the implementation of the Driver Diagnostics Protocol to see if a requested language is supported.

The **`RunDiagnostics()`** service runs diagnostics on the controller that a driver is managing or a child that the driver has produced. This service is not allowed to use any of the console-I/O-related protocols. Instead, the results of the diagnostics are returned to the caller in a buffer, and the caller may choose to log the results or display the results of the diagnostics that were executed. The format of the results must be clear and intuitive to the user.

The implementations of the Driver Diagnostics Protocol will change in complexity depending on the driver type. A device driver is fairly simple to implement. A bus driver or a hybrid driver may be more complex because it may provide diagnostics for both the bus controller and the child controllers. These implementations will be discussed in the following sections.

The **EFI_DRIVER_DIAGNOSTICS_PROTOCOL** must be installed on the same handle as the **EFI_DRIVER_BINDING_PROTOCOL**. This install operation is done in the driver's entry point. Installing the **EFI_DRIVER_BINDING_PROTOCOL** and the **EFI_DRIVER_DIAGNOSTICS_PROTOCOL** is covered in section 7.1. If an error is generated when installing the Driver Diagnostics Protocol, then this error should not cause the entire driver to fail.

The implementation of the Driver Diagnostics Protocol for a specific driver is typically found in the file **DriverDiagnostics.c**. This file contains the instance of the **EFI_DRIVER_DIAGNOSTIOCS_PROTOCOL** along with the implementation of **RunDiagnostics()**. Example 12-1 below shows the template for the Driver Diagnostics Protocol.

```
EFI_DRIVER_DIAGNOSTICS_PROTOCOL g<<DriverName>>DriverDiagnostics = {
  <<DriverName>>DriverDiagnosticsRunDiagnostics,
  "eng"
};

EFI_STATUS
<<DriverName>>DriverDiagnosticsRunDiagnostics (
  IN  EFI_DRIVER_DIAGNOSTICS_PROTOCOL  *This,
  IN  EFI_HANDLE                       ControllerHandle,
  IN  EFI_HANDLE                       ChildHandle  OPTIONAL,
  IN  EFI_DRIVER_DIAGNOSTIC_TYPE       DiagnosticType,
  IN  CHAR8                            *Language,
  OUT EFI_GUID                         **ErrorType,
  OUT UINTN                            *BufferSize,
  OUT CHAR16                           **Buffer
  )
{
}
```

Example 12-1. Driver Diagnostics Protocol Template

The *DiagnosticType* parameter tells the driver the type of diagnostics to perform. Standard diagnostics must be implemented, and they test basic functionality and should complete in less than 30 seconds. Extended diagnostics are recommended and may take more than 30 seconds to execute. Manufacturing diagnostics are intended to be used in a manufacturing and test environment.

ErrorType, *BufferSize*, and *Buffer* are the return parameters that report the results of the diagnostic. *Buffer* begins with a **NULL**-terminated Unicode string, so the caller of the **RunDiagnostics()** service can display a human-readable diagnostic result. *ErrorType* is a GUID that defines the format of the data buffer that follows the **NULL**-terminated Unicode string. *BufferSize* is the size of *Buffer* that includes the **NULL**-terminated Unicode string and the GUID-specific data buffer. The implementation of **RunDiagnostics()** must allocate *Buffer* using the service **gBS->AllocatePool()**, and it is the caller's responsibility to free this buffer with **gBS->FreePool()**.

Once a Driver Diagnostics Protocol is implemented, it can be tested with the EFI Shell command **drvdiag**. Platform vendors are also expected to implement extensions to the EFI boot manager that allow the user to execute diagnostics on the various devices in a platform. These EFI boot manager extensions use the services of the Component Name Protocol to display the names of devices, and the services of the Driver Diagnostics Protocol are used to execute diagnostics on each device.

The platform vendor also has the ability to automatically execute diagnostics on devices each time the platform is booted. If all the diagnostics are executed, then the boot time may increase. If none of the diagnostics is executed, then there is a chance an operating system may fail to boot due to a failure that could have been detected. The platform vendor will have to make a policy decision regarding diagnostics and may choose to add setup options that allow the user to enable or disable the execution of diagnostics on each boot.

The Driver Diagnostics Protocol is available only for devices that a driver is currently managing. Because EFI 1.10 supports connecting the minimum number of drivers and devices that are required to establish console and gain access to the boot device, there may be many unconnected devices that support diagnostics. As a result, when the user wishes to enter a “platform configuration” mode, the EFI boot manager will be required to connect all drivers to all devices, so that the user will be able to see all the devices that support diagnostics.

12.1 Device Drivers

Device drivers that implement the Driver Diagnostics Protocols need to make sure that *ChildHandle* is **NULL** and that *ControllerHandle* represents a device that the driver is currently managing. Example 11-3 shows the steps required to check these parameters and retrieve the private context data structure. If these checks pass, then the diagnostic will be executed and results will be returned. The diagnostic code will typically use the services of the protocols that the driver produces and the services of the protocols that the driver consumes to verify the operation of the controller. For example, a PCI device driver that consumes the PCI I/O Protocol and produces the Block I/O Protocol can use the services of the PCI I/O Protocol to verify the operation of the PCI controller. The Block I/O Services can be used to verify that the entire driver is working as expected.

12.2 Bus Drivers and Hybrid Drivers

Bus drivers and hybrid drivers that implement the Driver Diagnostics Protocols need to make sure that *ControllerHandle* and *ChildHandle* represent a device that the driver is currently managing. Example 11-4 shows the steps that are required to check these parameters and retrieve the private context data structure. If these checks pass, then the diagnostic will be executed and the results will be returned. The diagnostic code will typically use the services of the protocols that the driver produces and the services of the protocols that the drivers consumes to verify the operation of the controller. For example, a PCI device driver that consumes the PCI I/O Protocol and produces the Block I/O Protocol can use the services of the PCI I/O Protocol to verify the operation of the PCI controller. The Block I/O Services can be used to verify that the entire driver is working as expected. Bus drivers and hybrid drivers should provide diagnostics for both the bus controller and the child controllers that these types of drivers produce. Implementing diagnostics for only the bus controller or only the child controllers is strongly discouraged.

12.3 Implementing RunDiagnostics() as an Application

One possible design of the Driver Diagnostics Protocol is to implement the diagnostics as an EFI application that is stored with the device or in an EFI system partition. The implementation of **RunDiagnostics()** would be changed so that it does not directly execute the diagnostics. It would likely perform the same parameter checks as before and it would still retrieve the private context data structure. Then, instead of executing diagnostics, it would use the **gBS->LoadImage()** and **gBS->StartImage()** services to load and execute the EFI application that runs the diagnostics. This application would then return the results of the diagnostics back to **RunDiagnostics()**.

13
Bus Specific Driver Override Protocol

Some bus drivers are required to produce the Bus Specific Driver Override Protocol. The driver model for a specific bus type must declare if this protocol is required or not. In general, this protocol applies only to bus types that provide containers for drivers on their child devices. At this time, the only bus type that is required to produce this protocol is PCI, and the container for drivers is the PCI option ROM. The PCI bus driver is required to produce the Bus Specific Driver Override Protocol for PCI devices that have an attached PCI option ROM such that the PCI option ROM contains one or more loadable EFI drivers. If a PCI option ROM is not present or the PCI option ROM does not contain any loadable EFI drivers, than a Bus Specific Driver Override Protocol will not be produced for that PCI device.

The Bus Specific Driver Override Protocol is consumed only by the EFI Boot Service **`ConnectController()`** to determine the order that EFI drivers are used to attempt to start a device. An EFI driver never consumes the Bus Specific Driver Override Protocol.

The Bus Specific Driver Override Protocol is simple to implement. It contains one service called **`GetDriver()`** that returns an ordered list of image handles for the EFI drivers that were loaded from the EFI driver container. For PCI, the order in which the image handles are returned matches the order in which the EFI drivers were found in the PCI option ROM, from the lowest address to the highest address. The PCI bus driver is responsible for enumerating the PCI devices on a PCI bus. When a PCI device is discovered, the PCI device is also checked to see if it has an attached PCI option ROM. The PCI option ROM contents must follow the *PCI 2.2 Specification* for storing one or more images. The PCI bus driver will walk the list of images in a PCI option ROM looking for EFI drivers. If an EFI driver is found, it is optionally decompressed using the Decompress Protocol and then loaded, and the driver entry point is called using the EFI Boot Services **`LoadImage()`** and **`StartImage()`**. If **`LoadImage()`** does not return an error, then the EFI driver must be added to the end of the list of drivers that the Bus Specific Driver Override Protocol for that PCI device will return when its **`GetDriver()`** service is called. This addition could be implemented in many ways, including an array of image handles or a linked list of image handles.

The following group of code examples shows an implementation of the Bus Specific Driver Override Protocol using an array allocated from the pool to manage the list of EFI driver image handles. Example 13-1 shows a fragment of the private context structure that is used to manage the child-device-related information in a bus driver. Most private data structures in EFI drivers should contain a *`Signature`* field. This field is used in debug builds to make sure that the correct structure is being used by an EFI driver. If a bad pointer is passed into a function that examines the signature, then an **`ASSERT()`** will be generated. The signature of “Priv” is just an example. Each EFI driver should create a new signature value. The *`BusSpecificDriverOverride`* field is the protocol instance for the Bus Specific Driver Override Protocol. The *`NumberOfHandles`* field is the number of image handles that the **`GetDriver()`** function of the Bus Specific Driver Override Protocol can return. The *`HandleBufferSize`* field is the number of handles that can be stored in the array *`HandleBuffer`*, and the *`HandleBuffer`* field is the array of image handles that may be returned by the **`GetDriver()`** function of the Bus Specific Driver Override

Protocol. The **CR()** macro at the bottom of Example 13-1 will convert the *This* pointer for the Bus Specific Driver Override Protocol to a pointer to the **PRIVATE** structure. This pointer is used by the **GetDriver()** function to retrieve the private context structure.

```
#define PRIVATE_DATA_SIGNATURE   EFI_SIGNATURE_32('P','r','i','v')

typedef struct {
  UINTN                                          Signature;
  . . .
  EFI_PCI_BUS_SPECIFIC_DRIVER_OVERRIDE_PROTOCOL  BusSpecificDriverOverride;
  UINTN                                          NumberOfHandles;
  UINTN                                          HandleBufferSize;
  EFI_HANDLE                                     *HandleBuffer;
  . . .
} PRIVATE;

#define PRIVATE_FROM_BUS_SPECIFIC_DRIVER_OBVERRIDE_THIS(a) \
  CR(a, PRIVATE, BusSpecificDriverOverride, PRIVATE_DATA_SIGNATURE)
```

Example 13-1. Private Context for a Bus Specific Driver Override Protocol

Example 13-2 is an example implementation of the **GetDriver()** function of the Bus Specific Driver Override Protocol. The first step is to retrieve the private context structure from the *This* pointer. This retrieval is done with the **CR()** macro from Example 13-1. The next step is to see if there are any image handles in the private structure. If there are none, then **EFI_NOT_FOUND** is returned. The next step is to see if *DriverImageHandle* is a pointer to **NULL**. If it is, then the first image handle from *HandleBuffer* is returned. If *DriverImageHandle* is not a pointer to **NULL**, then a search is made through *HandleBuffer* to find a matching handle. If a matching handle is not found, then **EFI_INVALID_PARAMETER** is returned. If a matching handle is found, then the next handle in the array is returned. If the matching handle is the last handle in the array, then **EFI_NOT_FOUND** is returned.

```
EFI_STATUS
GetDriver (
  IN     EFI_BUS_SPECIFIC_DRIVER_OVERRIDE_PROTOCOL  *This,
  IN OUT EFI_HANDLE                                 *DriverImageHandle
  )

{
  UINTN    Index;
  PRIVATE  Private;

  Private = PRIVATE_FROM_BUS_SPECIFIC_DRIVER_OVERRIDE_THIS(This);

  if (Private->NumberOfHandles == 0) {
    return EFI_NOT_FOUND;
  }
  if (*DriverImageHandle == NULL) {
    *DriverImageHandle = Private->HandleBuffer[0];
    return EFI_SUCCESS;
  }
  for (Index = 0; Index < Private->NumberOfHandles; Index++) {
    if (*DriverImageHandle == Private->HandleBuffer[Index]) {
      Index = Index + 1;
      if (Index < Private->NumberOfHandles) {
        *DriverImageHandle = Private->HandleBuffer[Index];
        return EFI_SUCCESS;
```

```
      } else {
        return EFI_NOT_FOUND;
      }
    }
  }
  return EFI_INVALID_PARAMETER;
}
```

Example 13-2. GetDriver() Function of a Bus Specific Driver Override Protocol

Example 13-3 is an example implementation of a worker function that adds the image handle of a driver to the array. This function would be used when the PCI bus driver is scanning the PCI option ROMs for EFI drivers. As each EFI driver is loaded, this function would be called to add the image handle of the EFI driver to the Bus Specific Driver Override Protocol for the PCI controller that is associated with the PCI option ROM.

If there is not enough room in the image handle array, then an array with ten more handles is allocated, the contents of the old array are transferred to the new array, and the old array is freed. If there is not enough memory to allocate the new array, then **EFI_OUT_OF_RESOURCES** is returned. Once there is enough room to store the new image handle, the image handle is added to the end of the array and **EFI_SUCCESS** is returned.

```
EFI_STATUS
AddDriver(
  IN PRIVATE     *Private,
  IN EFI_HANDLE  DriverImageHandle
  )

{
  EFI_STATUS  Status;
  EFI_HANDLE  *NewBuffer;

  if (Private->NumberOfHandles >= Private->HandleBufferSize) {
    Status = gBS->AllocatePool (
                    EfiBootServicesMemory,
                    (Private->HandleBufferSize + 10) * sizeof (EFI_HANDLE),
                    (VOID **)&NewBuffer
                    );
    if (EFI_ERROR (Status)) {
      return Status;
    }

    gBS->CopyMem (
           NewBuffer,
           Private->HandleBuffer,
           Private->HandleBufferSize * sizeof (EFI_HANDLE)
           );

    Private->HandleBufferSize += 10;
    if (Private->HandleBuffer) {
      gBS->FreePool ( Private->HandleBuffer);
    }
    Private->HandleBuffer = NewBuffer;
  }

  Private->HandleBuffer[Private->NumberOfHandles] = DriverImageHandle;
```

```
  Private->NumberOfHandles++;
  return EFI_SUCCESS;
}
```

Example 13-3. AddDriver() Worker Function

Example 13-4 is an example implementation of a constructor function that initializes the fields of the private context structure that are related to this implementation of the Bus Specific Driver Override Protocol. There is only one service in the Bus Specific Driver Override Protocol, so the **GetDriver()** service is initialized. In addition, the array of image handles is initialized to contain zero handles. The first call to the **AddDriver()** worker function will allocate space for up to ten driver image handles.

```
EFI_STATUS
InitializeBusSpecificDriverOverrideInstance (
  PRIVATE  *Private
  )

{
  Private->BusSpecificDriverOverride.GetDriver = GetDriver;
  Private->NumberOfHandles = 0;
  Private->HandleBuffer = NULL;
  Private->HandleBufferSize = 0;
  return EFI_SUCCESS;
}
```

Example 13-4. Bus Specific Driver Override Protocol Constructor

14
PCI Driver Design Guidelines

There are several classes of PCI drivers that cooperate to provide support for PCI controllers in a platform. Table 14-1 lists these PCI drivers.

Table 14-1. Classes of PCI Drivers

Class of Driver	Description
Root bridge driver	Produces one or more instances of the PCI Root Bridge I/O Protocol.
PCI bus driver	Consumes the PCI Root Bridge I/O Protocol, produces a child handle for each PCI controller, and installs the Device Path Protocol and the PCI I/O Protocol onto each child handle.
PCI driver	Consumes the PCI I/O Protocol and produces an I/O abstraction that provides services for the console devices and boot devices that are required to boot an EFI-compliant operating system.

This chapter will concentrate on the design and implementation of PCI drivers. PCI drivers must follow all of the general design guidelines described in chapter 5. In addition, this chapter covers the guidelines that apply specifically to the management of PCI controllers.

Figure 14-1 below shows an example PCI driver stack and the protocols that the PCI-related drivers consume and produce. In this example, the platform hardware produces a single PCI root bridge. The PCI Root Bridge I/O Protocol driver accesses the hardware resources to produce a single handle with the **EFI_DEVICE_PATH_PROTOCOL** and the **EFI_PCI_ROOT_BRIDGE_IO_PROTOCOL**. The PCI bus driver consumes the services of the **PCI_ROOT_BRIDGE_IO_PROTOCOL**, and uses those services to enumerate the PCI controllers present in the system. In this example, the PCI bus driver detected a disk controller, a graphics controller, and a USB host controller. As a result, the PCI bus driver produces three child handles with the **EFI_DEVICE_PATH_PROTOCOL** and the **EFI_PCI_IO_PROTOCOL**. The driver for the PCI disk controller consumes the services of the **EFI_PCI_IO_PROTOCOL** and produces two child handles with the **EFI_DEVICE_PATH_PROTOCOL** and the **EFI_BLOCK_IO_PROTOCOL**. The PCI driver for the graphics controller consumes the services of the **EFI_PCI_IO_PROTOCOL** and produces the **EFI_UGA_DRAW_PROTOCOL**. Finally, the PCI driver for the USB host controller consumes the services of the **EFI_PCI_IO_PROTOCOL** to produce the **EFI_USB_HOST_CONTROLLER_PROTOCOL**. It is not shown in Figure 14-1, but the **EFI_USB_HOST_CONTROLLER_PROTOCOL** would then be consumed by the USB bus driver to produce child handles for each USB device, and USB drivers would then manage those child handles. Chapter 15 contains the guidelines for designing USB drivers.

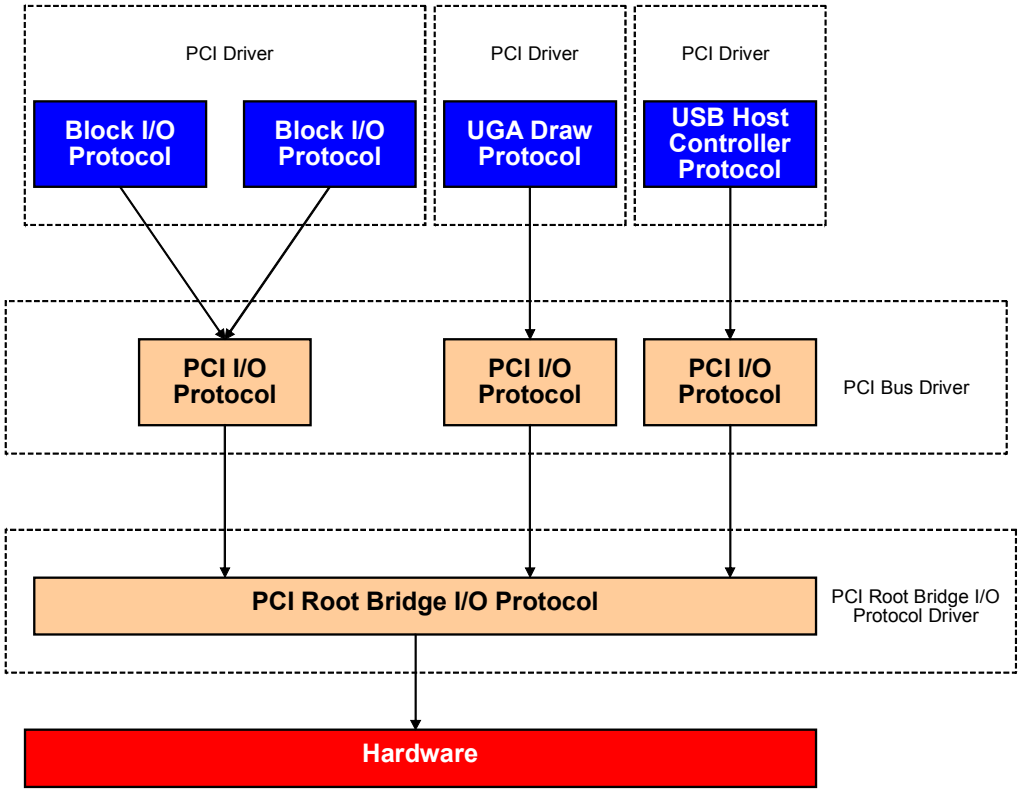

Figure 14-1. PCI Driver Stack

14.1 PCI Root Bridge I/O Protocol Drivers

An OEM or IBV typically implements the root bridge driver that produces the PCI Root Bridge I/O Protocol. This code is chipset specific and directly accesses the chipset resources to produce the services of the PCI Root Bridge I/O Protocol. A sample driver for systems with a PC-AT-compatible chipset is included in the *EFI 1.10 Sample Implementation*. The source code to this driver is in the directory **`\Efi1.1\Edk\Drivers\PcatPciRootBridge`**.

14.2 PCI Bus Drivers

The *EFI 1.10 Sample Implementation* contains a generic PCI bus driver. This driver uses the services of the **`EFI_PCI_ROOT_BRIDGE_IO_PROTOCOL`** to enumerate PCI devices and produce child handle with an **`EFI_DEVICE_PATH_PROTOCOL`** and an **`EFI_PCI_IO_PROTOCOL`**. This bus type can support producing one child handle at a time by parsing the *`RemainingDevicePath`* in its **`Supported()`** and **`Start()`** services. However, producing one child handle at a time generally does not make sense because the PCI bus driver needs to enumerate and assign resources to all of the PCI devices before even a single child handle

can be produced. It does not take much extra time to produce the child handles for all the PCI devices that were enumerated, so it is recommended that the PCI bus driver produce all of the PCI devices on the first call to **Start()**.

If EFI-based system firmware is ported to a new platform, most of the PCI-related changes occur in the implementation of the root bridge driver that produces the **EFI_PCI_ROOT_BRIDGE_IO_PROTOCOL**. Customization of the PCI bus driver is discouraged. As a result, the design and implementation of the PCI bus driver will not be covered in detail.

14.2.1 Hot-Plug PCI Buses

The PCI bus driver in the *EFI 1.10 Sample Implementation* does not support hot-plug events in the preboot environment. The PCI bus driver will function correctly with hot-plug-capable hardware, but the hot-add, hot-remove, and hot-replace events will be allowed only while an OS that supports hot-plug events is executing. The PCI bus driver would require updates to support hot-plug events in the preboot environment.

14.3 PCI Drivers

PCI drivers use the services of the **EFI_PCI_IO_PROTOCOL** to produce one or more protocols that provide I/O abstractions for a PCI controller. PCI drivers follow the EFI Driver Model, so they may be any of the following:

- Device drivers
- Bus drivers
- Hybrid drivers

The PCI drivers for graphics controllers are typically device drivers that consume the **EFI_PCI_IO_PROTOCOL** and produce the **EFI_UGA_DRAW_PROTOCOL** and **EFI_UGA_IO_PROTOCOL**. The PCI drivers for USB host controllers are typically device drivers that consume the **EFI_PCI_IO_PROTOCOL** and produce the **EFI_USB_HOST_CONTROLLER_PROTOCOL**. The PCI drivers for disk controllers are typically bus drivers or hybrid drivers that consume the **EFI_PCI_IO_PROTOCOL** and **EFI_DEVICE_PATH_PROTOCOL** and produce child handles with the **EFI_DEVICE_PATH_PROTOCOL** and **EFI_BLOCK_IO_PROTOCOL**. PCI drivers for disk controllers that use the SCSI command set will also typically produce the **EFI_SCSI_PASS_THRU_PROTOCOL** for each SCSI channel that the disk controller produces. Chapter 16 covers the details on SCSI drivers.

14.3.1 Supported()

A PCI driver must implement the **EFI_DRIVER_BINDING_PROTOCOL** that contains the **Supported()**, **Start()**, and **Stop()** services. The **Supported()** service evaluates the *ControllerHandle* that is passed in to see if the *ControllerHandle* represents a PCI device that the PCI driver knows how to manage. The most common method of implementing this test is for the PCI driver to retrieve the PCI configuration header from the PCI controller and evaluate the device ID, vendor ID, and possibly the class code fields of the PCI configuration header. If these fields match the values that the PCI driver knows how to manage, then

Supported() returns **EFI_SUCCESS**. Otherwise, the **Supported()** service will return **EFI_UNSUPPORTED**. The PCI driver must be careful not to disturb the state of the PCI controller because a different PCI driver may currently manage the PCI controller.

Example 14-1 below shows an example of the **Supported()** service for the XYZ PCI driver that manages a PCI controller with a vendor ID of 0x8086 and a device ID of 0xFFFE. First, it attempts to open the PCI I/O Protocol **BY_DRIVER** with **OpenProtocol()**. If the PCI I/O Protocol cannot be opened, then the PCI driver does not support the controller specified by *ControllerHandle*. If the PCI I/O Protocol is opened, then the services of the PCI I/O Protocol are used to read the vendor ID and device ID from the PCI configuration header. The PCI I/O Protocol is always closed with **CloseProtocol()**, and **EFI_SUCCESS** is returned if the vendor ID and device ID match.

```
#define XYZ_VENDOR_ID  0x8086
#define XYZ_DEVICE_ID  0xFFFE

EFI_STATUS
EFIAPI
XyzDriverBindingSupported (
  IN EFI_DRIVER_BINDING_PROTOCOL  *This,
  IN EFI_HANDLE                   ControllerHandle,
  IN EFI_DEVICE_PATH_PROTOCOL     *RemainingDevicePath  OPTIONAL
  )
{
  EFI_STATUS           Status;
  EFI_PCI_IO_PROTOCOL  *PciIo;
  UINT16               VendorId;
  UINT16               DeviceId;

  //
  // Open the PCI I/O Protocol on ControllerHandle
  //
  Status = gBS->OpenProtocol (
                 ControllerHandle,
                 &gEfiPciIoProtocolGuid,
                 (VOID **) &PciIo,
                 This->DriverBindingHandle,
                 ControllerHandle,
                 EFI_OPEN_PROTOCOL_BY_DRIVER
                 );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Read the vendor ID from the PCI configuration header
  //
  Status = PciIo->Pci.Read (
                       PciIo,
                       EfiPciIoWidthUint16,
                       0,
                       sizeof (VendorId),
                       &VendorId
                       );
  if (EFI_ERROR (Status)) {
    goto Done;
  }
```

```
  //
  // Read the device ID from the PCI configuration header
  //
  Status = PciIo->Pci.Read (
                        PciIo,
                        EfiPciIoWidthUint16,
                        2,
                        sizeof (DeviceId),
                        &DeviceId
                        );
  if (EFI_ERROR (Status)) {
    goto Done;
  }

  //
  // Evaluate VendorId and DeviceId
  //
  Status = EFI_SUCCESS;
  if (VendorId != XYZ_VENDOR_ID || DeviceId != XYZ_DEVICE_ID) {
    Status = EFI_UNSUPPORTED;
  }

Done:
  //
  // Close the PCI I/O Protocol
  //

  gBS->CloseProtocol (
         ControllerHandle,
         &gEfiPciIoProtocolGuid,
         This->DriverBindingHandle,
         ControllerHandle
         );

  return Status;
}
```

Example 14-1. Supported() Service with Partial PCI Configuration Header

The previous example performs two 16-bit reads from the PCI configuration header. The code would be smaller if the entire PCI configuration header was read at once. However, this would increase the execution time because the **`Supported()`** service would read the entire PCI configuration header for every *`ControllerHandle`* that was passed in. The **`Supported()`** service is supposed to be a small, quick check. If a more extensive evaluation of the PCI configuration header is required, then it may make sense to read the entire PCI configuration header at once. Example 14-2 below shows the same example as above, except it reads the entire PCI configuration header at once in 32-bit chunks.

```
#define XYZ_VENDOR_ID  0x8086
#define XYZ_DEVICE_ID  0xFFFE

EFI_STATUS
EFIAPI
XyzDriverBindingSupported (
  IN EFI_DRIVER_BINDING_PROTOCOL  *This,
  IN EFI_HANDLE                   ControllerHandle,
  IN EFI_DEVICE_PATH_PROTOCOL     *RemainingDevicePath  OPTIONAL
```

```
  )
{
  EFI_STATUS           Status;
  EFI_PCI_IO_PROTOCOL  *PciIo;
  PCI_TYPE00           Pci;

  //
  // Open the PCI I/O Protocol on ControllerHandle
  //
  Status = gBS->OpenProtocol (
                 ControllerHandle,
                 &gEfiPciIoProtocolGuid,
                 (VOID **) &PciIo,
                 This->DriverBindingHandle,
                 ControllerHandle,
                 EFI_OPEN_PROTOCOL_BY_DRIVER
                 );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Read the entire PCI configuration header
  //
  Status = PciIo->Pci.Read (
                       PciIo,
                       EfiPciIoWidthUint32,
                       0,
                       sizeof (Pci) / sizeof (UINT32),
                       &Pci
                       );
  if (EFI_ERROR (Status)) {
    goto Done;
  }

  //
  // Evaluate VendorId and DeviceId
  //
  Status = EFI_SUCCESS;
  if (Pci.Header.VendorId != XYZ_VENDOR_ID ||
      Pci.Header.DeviceId != XYZ_DEVICE_ID    ) {
    Status = EFI_UNSUPPORTED;
  }

Done:
  //
  // Close the PCI I/O Protocol
  //

  gBS->CloseProtocol (
         ControllerHandle,
         &gEfiPciIoProtocolGuid,
         This->DriverBindingHandle,
         ControllerHandle
         );

  return Status;
}
```

Example 14-2. Supported() Service with Entire PCI Configuration Header

14.3.2 Start() and Stop()

The **Start()** service of the Driver Binding Protocol for a PCI driver also opens the PCI I/O Protocol **BY_DRIVER**. If the PCI driver is a bus driver or a hybrid driver, then the Device Path Protocol will also be opened **BY_DRIVER**. In addition, all PCI drivers are required to call the **Attributes()** service of the PCI I/O Protocol to enable the I/O, memory, and bus master bits in the Command register of the PCI configuration header. By default, the PCI bus driver is not required to enable the Command register of the PCI controllers. Instead, it is the responsibility of the **Start()** service to enable these bits and of the **Stop()** service to disable these bits.

There is one additional attribute that must be specified in this call to the **Attributes()** service. If the PCI controller is a bus master and capable of generating 64-bit DMA addresses, then the **EFI_PCI_IO_ATTRIBUTE_DUAL_ADDRESS_CYCLE** attribute must also be enabled. Unfortunately, there is no standard method for detecting if a PCI controller supports 32-bit or 64-bit DMA addresses. As a result, it is the PCI driver's responsibility to inform the PCI bus driver that the PCI controller is capable of producing 64-bit DMA addresses. The PCI bus driver will assume that all PCI controllers are only capable of generating 32-bit DMA addresses unless the PCI driver enables the dual address cycle attribute. The PCI bus driver uses this information along with the services of the PCI Root Bridge I/O Protocol to optimize PCI DMA transactions. If a PCI bus master that is capable of 32-bit DMA addresses is present in a platform that supports more than 4 GB of system memory, then the DMA transactions may have to be double buffered, and this double buffering can reduce the performance of a driver. It is also possible for some platforms to only support system memory above 4 GB. For these reasons, a PCI driver must always accurately describe the DMA capabilities of the PCI controller from the **Start()** service of the Driver Binding Protocol.

Example 14-3 below shows the code fragment from the **Start()** and **Stop()** services of a PCI driver for a PCI controller that supports 64-bit DMA transactions.

```
EFI_STATUS           Status;
EFI_PCI_IO_PROTOCOL  *PciIo;

//
// Attributes() call from Start() service after the PCI I/O Protocol is
// opened
//
Status = PciIo->Attributes (
                  PciIo,
                  EfiPciIoAttributeOperationEnable,
                  EFI_PCI_DEVICE_ENABLE |
                  EFI_PCI_IO_ATTRIBUTE_DUAL_ADDRESS_CYCLE,
                  NULL
                  );

//
// Attributes() call from Stop() service before the PCI I/O Protocol is
// closed
//
Status = PciIo->Attributes (
                  PciIo,
                  EfiPciIoAttributeOperationDisable,
                  EFI_PCI_DEVICE_ENABLE |
```

```
                    EFI_PCI_IO_ATTRIBUTE_DUAL_ADDRESS_CYCLE,
                    NULL
                    );
```

Example 14-3. Start() and Stop() for a 64-Bit DMA Capable PCI Controller

Example 14-4 below shows the code fragment from the **Start()** and **Stop()** services of a PCI driver for a PCI controller that does not support 64-bit DMA transactions.

```
EFI_STATUS            Status;
EFI_PCI_IO_PROTOCOL   *PciIo;

//
// Attributes() call from Start() service after the PCI I/O Protocol is
// opened
//
Status = PciIo->Attributes (
                    PciIo,
                    EfiPciIoAttributeOperationEnable,
                    EFI_PCI_DEVICE_ENABLE,
                    NULL
                    );

//
// Attributes() call from Stop() service before the PCI I/O Protocol is
// closed
//
Status = PciIo->Attributes (
                    PciIo,
                    EfiPciIoAttributeOperationDisable,
                    EFI_PCI_DEVICE_ENABLE,
                    NULL
                    );
```

Example 14-4. Start() and Stop() for a 32-Bit DMA Capable PCI Controller

Table 14-2 lists the **#define** statements that can be used with the **Attributes()** service. A PCI driver must use the **Attributes()** service to enable the decodes on the PCI controller, accurately describe the PCI controller DMA capabilities, and request that specific ISA addresses be forwarded to the PCI controller if the PCI controller requires ISA resources. The call to **Attributes()** will fail if the request cannot be satisfied, and if this failure occurs, then the **Start()** function should return an error. Once again, any attributes that are enabled in the **Start()** service must be disabled in the **Stop()** service.

Table 14-2. PCI Attributes

Attribute	Description
EFI_PCI_IO_ATTRIBUTE_ISA_MOTHERBOARD_IO	Used to request the forwarding of I/O cycles 0x0000–0x00FF (10-bit decode).
EFI_PCI_IO_ATTRIBUTE_ISA_IO	Used to request the forwarding of I/O cycles 0x0000–0x0FFF (16-bit decode).
EFI_PCI_IO_ATTRIBUTE_VGA_PALETTE_IO	Used to request the forwarding of I/O cycles 0x3C6, 0x3C8, and 0x3C9 (10-bit decode).

continued

Table 14-2. PCI Attributes (continued)

Attribute	Description
EFI_PCI_IO_ATTRIBUTE_VGA_MEMORY	Used to request the forwarding of MMIO cycles 0xA0000–0xBFFFF (24-bit decode).
EFI_PCI_IO_ATTRIBUTE_VGA_IO	Used to request the forwarding of I/O cycles 0x3B0–0x3BB and 0x3C0–0x3DF (10-bit decode).
EFI_PCI_IO_ATTRIBUTE_IDE_PRIMARY_IO	Used to request the forwarding of I/O cycles 0x1F0–0x1F7, 0x3F6, 0x3F7 (10-bit decode).
EFI_PCI_IO_ATTRIBUTE_IDE_SECONDARY_IO	Used to request the forwarding of I/O cycles 0x170–0x177, 0x376, 0x377 (10-bit decode).
EFI_PCI_IO_ATTRIBUTE_IO	Enable the I/O decode bit in the Command register.
EFI_PCI_IO_ATTRIBUTE_MEMORY	Enable the Memory decode bit in the Command register.
EFI_PCI_IO_ATTRIBUTE_BUS_MASTER	Enable the Bus Master bit in the Command register.
EFI_PCI_IO_ATTRIBUTE_DUAL_ADDRESS_CYCLE	Clear for PCI controllers that cannot generate a DAC.
EFI_PCI_DEVICE_ENABLE	Enable I/O, Memory, and Bus Master bits in the Command register.

Table 14-3 lists the **`#define`** statements that can be used with the **`GetBarAttributes()`** and **`SetBarAttributes()`** services to adjust the attributes of a memory-mapped I/O region that is described by a Base Address Register (BAR) of a PCI controller. The support of these attributes is optional, and in general, a PCI driver uses these attributes to provide hints that may be used to improve the performance of a PCI driver. This improved performance is especially important for PCI drivers that manage graphics controllers. Any BAR attributes that are set in the **`Start()`** service must be cleared in the **`Stop()`** service.

Table 14-3. PCI BAR Attributes

Attribute	Description
EFI_PCI_IO_ATTRIBUTE_MEMORY_WRITE_COMBINE	Used to map a memory range of a BAR so writes are combined.
EFI_PCI_IO_ATTRIBUTE_MEMORY_CACHED	Used to map a memory range of a BAR so all read-write accesses are cached.
EFI_PCI_IO_ATTRIBUTE_MEMORY_DISABLE	Used to disable a memory range of a BAR.

Table 14-4 lists the **`#define`** statements that describe some additional attributes of a PCI controller. A PCI driver may retrieve the attributes of a PCI controller with the **`Attributes()`** service and check to see if these bits are set. The PCI driver may contain different code paths for embedded PCI controllers and PCI controllers that are present on add-in adapters.

Table 14-4. PCI Embedded Device Attributes

Attribute	Description
EFI_PCI_IO_ATTRIBUTE_EMBEDDED_DEVICE	Clear for an add-in PCI device.
EFI_PCI_IO_ATTRIBUTE_EMBEDDED_ROM	Clear for a physical PCI option ROM accessed through a ROM BAR.

14.4 Accessing PCI Resources

PCI drivers should only access the I/O and memory-mapped I/O resources on the PCI controllers that they manage. They should never attempt to access the I/O or memory-mapped I/O resource of a PCI controller that they are not managing. They should also never touch the I/O or memory-mapped I/O resources of the chipset or the motherboard.

The PCI I/O Protocol provides services that allow a PCI driver to easily access the resources of the PCI controllers that it is currently managing. These services hide platform-specific implementation details and prevent a PCI driver from inadvertently accessing the resources of the motherboard or other PCI controllers. The PCI I/O Protocol has also been designed to simplify the implementation of PCI drivers. For example, a PCI driver should never read the BARs in the PCI configuration header. Instead, the PCI driver passes in a `BarIndex` and `Offset` into the PCI I/O Protocol services. The PCI bus driver is responsible for managing the PCI controller's BARs.

The services of the PCI I/O Protocol that allow a PCI driver to access the resources on a PCI controller include the following. Chapter 17 contains several examples on how these services should be used.

- **`PciIo->PollMem()`**
- **`PciIo->PollIo()`**
- **`PciIo->Mem.Read()`**
- **`PciIo->Mem.Write()`**
- **`PciIo->Io.Read()`**
- **`PciIo->Io.Write()`**
- **`PciIo->Pci.Read()`**
- **`PciIo->Pci.Write()`**
- **`PciIo->CopyMem()`**

Another important resource that is provided through the PCI I/O Protocol is the PCI option ROM contents. The `RomSize` and `RomImage` fields of the PCI I/O Protocol provide access to a copy of the PCI option ROM contents. These fields may be useful if the PCI driver requires additional information from the contents of the PCI option ROM. It is important to note that the PCI option ROM contents cannot be modified through the `RomImage` field. Modifications to this buffer will only modify the copy of the PCI option ROM contents that are in system memory. The PCI I/O Protocol does not provide services to modify the contents of the actual PCI option ROM.

14.4.1 Memory-Mapped I/O Ordering Issues

PCI transactions follow the ordering rules defined in Appendix E of the *PCI 2.3 Specification.* The ordering rules vary for I/O, memory-mapped I/O, and PCI configuration cycles.

The PCI I/O Protocol **Mem.Read()** service generates PCI memory read cycles that are guaranteed to complete before control is returned to the PCI driver. However, the PCI I/O Protocol **Mem.Write()** service does not guarantee that the PCI memory cycles that are produced by this service will complete before control is returned to the PCI driver. This distinction means that memory write transactions may be sitting in write buffers when this service returns. If the PCI driver requires a **Mem.Write()** transaction to complete, then the **Mem.Write()** transaction must be followed by a **Mem.Read()** transaction to the same PCI controller. Some chipsets and PCI-to-PCI bridges are more sensitive to this issue than others. Example 14-5 below shows a **Mem.Write()** call to a memory-mapped I/O register at offset 0x20 into BAR #1 of a PCI controller. This write transaction is followed by a **Mem.Read()** call from the same memory-mapped I/O register. This combination will guarantee that the write transaction will be completed by the time the **Mem.Read()** call returns.

In general, this mechanism is not required because a PCI driver will typically read a status register, and this read transaction would force all posted write transactions to complete on the PCI controller. The only time this mechanism should be used is when a PCI driver performs a write transaction that is not followed by a read transaction and the PCI driver needs to guarantee that the write transaction is completed immediately.

```
EFI_PCI_IO_PROTOCOL  *PciIo;
UINT32               DmaStartAddress;
UINT16               Word;

//
// Write the value in DmaStartAddress to offset 0x20 of BAR #1
//
PciIo->Mem.Write (
             PciIo,
             EfiPciIoWidthUint32,
             1,
             0x20,
             1,
             &DmaStartAddress
             );

//
// Read offset 0x20 of BAR #1. This guarantees that the previous write
// transaction is posted to the PCI controller.
//
PciIo->Mem.Read (
             PciIo,
             EfiPciIoWidthUint32,
             1,
             0x20,
             1,
             &DmaStartAddress
             );
```

Example 14-5. Completing a Memory Write Transaction

14.4.2 Hardfail / Softfail

PCI drivers must make sure they do not access resources that are not allocated to any PCI controllers. Doing so may produce unpredictable results including platform hang conditions. For example, if a VGA device is in monochrome mode, accessing the VGA device's color registers may cause unpredictable results. The best rule of thumb here is to access only I/O or memory-mapped I/O resources to which the PCI driver knows for sure the PCI controller will respond. In general, this is not a concern because the PCI I/O Protocol services discussed in section 14.4 do not allow the PCI driver to access resources outside the resource ranges described in the BARs of the PCI controllers. However, two mechanisms allow a PCI driver to bypass these safeguards. The first is to use the **EFI_PCI_IO_PASS_THROUGH_BAR** with the PCI I/O Protocol services that provide access to I/O and memory-mapped I/O regions. The second is for a PCI driver to retrieve and use the services of a PCI Root Bridge I/O Protocol.

A PCI driver uses the **EFI_PCI_IO_PASS_THROUGH_BAR** to access ISA resources on a PCI controller. For a PCI driver to use this mechanism safely, the PCI driver must know that the PCI controller it is accessing will actually respond to the I/O or memory-mapped I/O requests in the ISA ranges. The PCI driver can typically know if it will respond by examining the class code, vendor ID, and device ID fields of the PCI controller in the PCI configuration header. The PCI driver must examine the PCI configuration header before any I/O or memory-mapped I/O operations are generated. The PCI configuration header is typically examined in the **Supported()** service, so it is safe to access the ISA resources in the **Start()** service and in the services of the I/O abstraction that the PCI driver is producing. Example 14-6 below shows an example using the **EFI_PCI_IO_PASS_THROUGH_BAR**.

```
EFI_PCI_IO_PROTOCOL  *PciIo;
UINT8                Data;
UINT16               Word;

//
// Write 0xAA to a Post Card at ISA address 0x80
//
Data = 0xAA;
PciIo->Io.Write(
            PciIo,
            EfiPciIoWidthUint8,
            EFI_PCI_IO_PASS_THROUGH_BAR,
            0x80,
            1,
            &Data
            );

//
// Read the first word from the VGA frame buffer
//
PciIo->Mem.Read(
             PciIo,
             EfiPciIoWidthUint16,
             EFI_PCI_IO_PASS_THROUGH_BAR,
             0xA0000,
             1,
             &Word
             );
```

Example 14-6. Accessing ISA Resources on a PCI Controller

A PCI driver must also take care when using the services of the PCI Root Bridge I/O Protocol. A PCI driver can retrieve the parent PCI Root Bridge I/O Protocol and use those services to touch any resource on the PCI bus. This touching can be very dangerous because the PCI driver may not know if a different PCI driver owns a resource or not. The use of this mechanism is strongly discouraged and will likely be used only by OEM drivers that have intimate knowledge of the platform and the chipset. Chapter 4 discusses the use of the **gBS->LocateDevicePath()** service, and the example associated with this service shows how the parent PCI Root Bridge I/O Protocol can be retrieved.

Instead of using the parent PCI Root Bridge I/O Protocol, PCI drivers that need to access the resources of other PCI controllers in the platform should search the handle database for controller handles that support the PCI I/O Protocol. To prevent resource conflicts, the PCI I/O Protocols from other PCI controllers should also be opened **BY_DRIVER**. Example 14-7 below shows how a PCI driver can easily retrieve the list of PCI controller handles in the handle database and use the services of the PCI I/O Protocol on each of those handles to find a peer PCI controller. For example, a PCI adapter that contains multiple PCI controllers behind a PCI-to-PCI bridge may use a single driver to manage all of the controllers on the adapter. When the PCI driver is connected to the first PCI controller on the adapter, the PCI driver will want to connect to all the other PCI controllers that have the same bus number as the first PCI controller. This example takes advantage of the **GetLocation()** service of the PCI I/O Protocol to find matching bus numbers.

```
EFI_STATUS            Status;
UINTN                 HandleCount;
EFI_HANDLE            *HandleBuffer;
UINTN                 Index;
EFI_PCI_IO_PROTOCOL   *PciIo;
UINTN                 MyBus;
UINTN                 Seg;
UINTN                 Bus;
UINTN                 Device;
UINTN                 Function;

//
// Retrieve the location of the PCI controller and store the bus number
// in MyBus.
//
Status = PciIo->GetLocation (PciIo, &Seg, &MyBus, &Device, &Function);
if (EFI_ERROR (Status)) {
  return Status;
}

//
// Retrieve the list of handles that support the PCI I/O protocol from
// the handle database. The number of handles that support the PCI I/O
// Protocol is returned in HandleCount, and the array of handle values is
// returned in HandleBuffer.
//
Status = gBS->LocateHandleBuffer (
                ByProtocol,
                &gEfiPciIoProtocolGuid,
                NULL,
                &HandleCount,
                &HandleBuffer
                );
if (EFI_ERROR (Status)) {
```

```
  return Status;
}

//
// Loop through all the handles the support the PCI I/O Protocol, and
// retrieve the instance of the PCI I/O Protocol. Use the BY_DRIVER
// open mode, so only PCI I/O Protocols that are not currently being
// managed will be considered.
//
for (Index = 0; Index < HandleCount; Index++) {
  Status = gBS->OpenProtocol (
                  HandleBuffer[Index],
                  &gEfiPciIoProtocolGuid,
                  (VOID **)&PciIo,
                  ImageHandle,
                  NULL,
                  EFI_OPEN_PROTOCOL_BY_DRIVER
                  );
  if (!EFI_ERROR (Status)) {
    //
    // Retrieve the location of the PCI controller and store the bus
    // number in Bus.
    //
    Status = PciIo->GetLocation (PciIo, &Seg, &Bus, &Device, &Function);
    if (!EFI_ERROR (Status)) {
      if (Bus == MyBus) {
        //
        // Store HandleBuffer[Index] so the driver knows it is managing
        // the PCI controller represented by HandleBuffer[Index]. This
        // would typically be stored in the private context data structure
        //

        //
        // Continue with the next PCI controller in the HandleBuffer array
        //
        continue;
      }
    }
  }

  //
  // Either the handle was already opened by another driver or the bus
  // numbers did not match, so close the PCI I/O Protocol and move
  // on to the next PCI handle.
  //
  gBS->CloseProtocol (
         HandleBuffer[Index],
         &gEfiPciIoProtocolGuid,
         ImageHandle,
         NULL,
         );
}

//
// Free the array of handles that was allocated by
// gBS->LocateHandleBuffer()
//
gBS->FreePool (HandleBuffer);
```

Example 14-7. Locate PCI Handles with Matching Bus Number

14.4.3 When a PCI Device Does Not Receive Resources

Some PCI controllers may require more resources than the PCI bus can offer. In such cases, the PCI controller must not be visible to PCI drivers because resources were not allocated to the PCI controller. The PCI bus driver should not create a child handle for a PCI controller that does not have allocated resources, and as a result, a PCI driver will never be passed a *ControllerHandle* for a PCI controller that does not have allocated resources. The platform vendor controls the policy decisions that are made when this type of resource-constrained condition is encountered. The PCI driver writer will never have to handle this case.

14.5 PCI DMA

There are three types of DMA transactions that can be implemented using the services of the PCI I/O Protocol:

- Bus master read transactions
- Bus master write transactions
- Common buffer transactions

The PCI I/O Protocol services that are used to manage PCI DMA transactions include the following:

- **PciIo->AllocateBuffer()**
- **PciIo->FreeBuffer()**
- **PciIo->Map()**
- **PciIo->Unmap()**
- **PciIo->Flush()**

14.5.1 Map() Service Cautions

One common mistake is in the use of the **Map()** service. The **Map()** service converts a system memory address to an address that can be used by the PCI device to perform a bus master DMA transaction. The device address that is returned is not related to the original system memory address. Some chipsets maintain a one-to-one mapping between system memory addresses and device addresses on the PCI bus. For this special case, the system memory address and device address will be the same. However, a PCI driver cannot tell if they are executing on a platform with this one-to-one mapping. As a result, a PCI driver must make as few assumptions about the system architecture as possible. Avoiding assumptions means that a PCI driver must never use the device address that is returned from **Map()** to access the contents of the DMA buffer. Instead, this value should only be used to program the base address of the DMA transaction into the PCI controller. This programming is typically accomplished with one or more I/O or memory-mapped I/O write transactions to the PCI controller that the PCI driver is managing.

Example 14-8 below shows the function prototype for the **Map()** service of the PCI I/O Protocol. A PCI driver can use *HostAddress* to access the contents of the DMA buffer, but the PCI driver should never use the returned parameter *DeviceAddress* to access the contents of the DMA buffer.

```
EFI_STATUS Map (
  IN     EFI_PCI_IO_PROTOCOL            *This,
  IN     EFI_PCI_IO_PROTOCOL_OPERATION  Operation,
  IN     VOID                           *HostAddress,
  IN OUT UINTN                          *NumberOfBytes,
  OUT    EFI_PHYSICAL_ADDRESS           *DeviceAddress,
  OUT    VOID                           **Mapping
  );
```

Example 14-8. Map() Function

14.5.2 Weakly Ordered Memory Transactions

Some processors, such as the Itanium processor, have weakly ordered memory models. This weak ordering means that system memory transactions may complete in a different order than the source code would seem to indicate. A PCI driver should be implemented so that the source code is compatible with as many processors and platforms as possible. As a result, the guidelines shown here should be followed even if the driver is not initially being written for an Itanium-based platform. The techniques shown here will not have any impact on the executable size of a driver for strongly ordered processors such as IA-32 and EBC.

14.5.3 Bus Master Read/Write Operations

When a DMA transaction is started or stopped, the ownership of the DMA buffer is transitioned from the processor to the DMA bus master and back to the processor. The PCI I/O Protocol provides the **Map()** and **Unmap()** services that are used to set up and complete a DMA transaction. The implementation of the PCI I/O Protocol in the PCI bus driver uses the **MEMORY_FENCE()** macro to guarantee that all system memory transactions from the processor are completed before the DMA transaction is started. This guarantee prevents the case where a DMA bus master reads from a location in the DMA buffer before a write transaction is flushed from the processor. Because this functionality is built into the PCI I/O Protocol itself, the PCI driver writer does not need to worry about this case for bus master read operations and bus master write operations.

A PCI driver is responsible for flushing all posted write data from a PCI controller when a bus master write operation is completed. First, the PCI driver should read from a register on the PCI controller to guarantee that all the posted write operations are flushed from the PCI controller and through any PCI-to-PCI bridges that are between the PCI controller and the PCI root bridge. Because PCI drivers are polled, they will typically read from a status register on the PCI controller to determine when the bus master write transaction is completed. This read operation is usually sufficient to flush the posted write buffers. The PCI driver must also call the **PciIo->Flush()** service at the end of a bus master write operation. This service flushes all the posted write buffers in the system chipset, so they are guaranteed to be committed to system memory. The combination of the read operation and the **PciIo->Flush()** call guarantee that the bus master's view of system memory and the processor's view of system memory are consistent. Example 14-9 below shows an example of how a bus master write transaction should be completed and guarantees that the bus master's view of system memory is consistent with the processor's view of system memory.

```
//
// Call PollIo() to poll for Bit #0 in register 0x24 of Bar #1 to be
// set to a 1. This example shows polling a status register to
// wait for a bus master write transaction to complete.
//
Status = PciIo->PollMem (
                  PciIo,
                  EfiPciIoWidthUint32,     // Width
                  1,                       // BarIndex
                  0x24,                    // Offset
                  0x01,                    // Mask
                  0x01,                    // Value
                  10000000,                // Poll for 1 second
                  &Result64                // Result
                  );
if (EFI_ERROR (Status)) {
  return Status;
}

Status = PciIo->Flush (PciIo);
if (EFI_ERROR (Status)) {
  return Status;
}
```

Example 14-9. Completing a Bus Master Write Operation

14.5.4 Bus Master Common Buffer Operations

Bus master common buffer operations are more complex to manage than bus master read operations and bus master write operations. This complexity is because both the bus master and the processor may simultaneously access a single region of system memory. The memory ordering of PCI transactions generated by the PCI bus master is defined in the *PCI Specification*. However, different processors may use different memory ordering models. As a result, common buffer operations should only be used when they are absolutely required.

If the common buffer memory region can be accessed in a single atomic processor transaction, then no hazards will be present. If the processor has deep write buffers, a write transaction can be delayed, so the **MEMORY_FENCE()** macro can be used to force all processor transactions to complete. If a memory region that the processor needs to read or write requires multiple atomic processor transactions, then hazards may exist if the operations are reordered. If the order that the processor transactions occur is important, then the **MEMORY_FENCE()** macro can be inserted between the processor transactions. However, inserting too many **MEMORY_FENCE()** macros will degrade system performance. For strongly ordered processors, the **MEMORY_FENCE()** macro is a no-op.

A good example where the **MEMORY_FENCE()** macro should be used is when a mailbox data structure is used to communicate between the processor and a bus master. The mailbox typically contains a valid bit that must be set by the processor after the processor has filled the contents of the mailbox. The bus master will scan the mailbox to see if the valid bit it set. When it sees the valid bit, it will read the rest of the mailbox contents and use them to perform an I/O operation. If the processor is weakly ordered, there is a chance that the valid bit will be set before the processor has written all of the other fields in the data structure. To resolve this issue, a **MEMORY_FENCE()** macro should be inserted just before and just after the valid bit is set.

Another mechanism that can be used to resolve these memory-ordering issues is the use of **VOLATILE** variables. If the data structure that is being used as a mailbox is declared in C as **VOLATILE**, then the C compiler will guarantee that all transactions to the **VOLATILE** data structure are strongly ordered. It is recommended that the **MEMORY_FENCE()** macro be used instead of **VOLATILE** data structures.

14.5.5 4 GB Memory Boundary

IA-32 platforms may support more than 4 GB of system memory, but EFI drivers for IA-32 platforms may only access memory below 4 GB. The 4 GB memory boundary becomes more complex on Itanium-based platforms. Itanium-based platforms do support more than 4 GB of system memory and EFI drivers running on these platforms can use the memory above and below 4 GB. It is important that EFI drivers are tested on Itanium-based platforms with both small and large memory configurations to make sure the EFI driver is not making any assumptions about the system memory configuration.

It is also important to note that some Itanium-based platforms may not map any system memory in the memory region below 4 GB. Instead, all the system memory is mapped above 4 GB. EFI drivers need to be designed to be compatible with these types of systems too. The general guideline for EFI drivers is to make as few assumptions about the memory configuration of the platform as possible. This guideline applies to the memory that an EFI driver allocates and the DMA buffers that a PCI bus master uses.

An EFI driver should not allocate buffers from specific addresses or below specific addresses. These types of allocations may fail on different system architectures. Likewise, the buffers used for DMA should not be allocated from a specific address or below a specific address. In addition, EFI drivers should always use the services of the PCI I/O Protocol to set up and complete DMA transactions. It is not legal to program a system memory address into a DMA bus master. This programming will work on chipsets that have a one-to-one mapping between system memory addresses and PCI DMA addresses, but it will not work with chipsets that remap DMA transactions. The sections that follow contain code examples for the different types of PCI DMA transactions that EFI supports and that show how the PCI I/O Protocol services should be used to maximize the platform compatibility of EFI drivers.

The *EFI Sample Implementation* contains an implementation of the PCI Root Bridge I/O Protocol for a PC-AT-compatible chipset that assumes a one-to-one mapping between system memory address and PCI DMA addresses. It also assumes that DMA operations are not supported above 4 GB. The implementation of the **Map()** and **Unmap()** services in the PCI Root Bridge I/O Protocol handle DMA requests above 4 GB by allocating a buffer below 4 GB and copying the data to the buffer below 4 GB. It is important to realize that these functions will be implemented differently for chipsets that do not assume a one-to-one mapping between system memory addresses and PCI DMA addresses.

14.5.6 DMA Bus Master Read Operation

The general algorithm for performing a bus master read operation is as follows:

- The processor initializes the contents of the DMA using *HostAddress*.
- Call **Map()** with an *Operation* of **EfiPciOperationBusMasterRead**.
- Program the DMA bus master with the *DeviceAddress* returned by **Map()**.
- Program the DMA bus master with the *NumberOfBytes* returned by **Map()**.
- Start the DMA bus master.
- Wait for DMA bus master to complete the bus master read operation.
- Call **Unmap()**.

The code example in Example 14-10 shows a function for performing a bus master read operation on a PCI controller. The PCI controller is accessed through the parameter *PciIo*. The system memory buffer that will be read by the bus master is specified by *HostAddress* and *Length*. This function performs one or more bus master read operations until either *Length* bytes have been read by the bus master or an error is detected. The PCI controller in this example has three MMIO registers in BAR #1. The MMIO register at offset 0x10 is a status register that the function uses to see if the DMA operation is complete or not. The function writes the start of the DMA transaction to the MMIO register at offset 0x20 and the length of the DMA transaction to the MMIO register at offset 0x24. The write operation to offset 0x24 also starts the DMA read operation. The services of the PCI I/O Protocol that are used in this example include **Map()**, **Unmap()**, **Mem.Write()**, and **PollMem()**. This example is for a 32-bit PCI bus master.

If a 64-bit PCI bus master was being used, then there would be two 32-bit MMIO registers to specify the start address and two 32-bit MMIO registers to specify the length. If the PCI bus master supports 64-bit DMA addressing, then the **EFI_PCI_ATTRIBUTE_DUAL_ADDRESS_CYCLE** attribute must be set in the **Start()** service of the PCI driver.

```
EFI_STATUS
DoBusMasterRead (
  IN EFI_PCI_IO_PROTOCOL  *PciIo,
  IN UINT8                *HostAddress,
  IN UINTN                *Length
  )

{
  EFI_STATUS            Status;
  UINTN                 NumberOfBytes;
  EFI_PHYSICAL_ADDRESS  DeviceAddress;
  VOID                  *Mapping;
  UINT32                DmaStartAddress;
  UINT32                ControllerStatus;

  //
  // Loop until the entire buffer specified by HostAddress and Length
  // has been read from the PCI DMA bus master
  //
  do {
    //
    // Call Map() to retrieve the DeviceAddress to use for the bus master
    // read operation. The Map() function may not support performing
    // a DMA operation for the entire length, so it may be broken up into
    // smaller DMA operations.
```

```
    //
    NumberOfBytes = *Length;
    Status = PciIo->Map (
                     PciIo,
                     EfiPciIoOperationBusMasterRead,
                     (VOID *)HostAddress,
                     &NumberOfBytes,
                     &DeviceAddress,
                     &Mapping
                     );
    if (EFI_ERROR (Status)) {
      return Status;
    }

    //
    // Write the DMA start address to MMIO Register 0x20 of Bar #1
    //
    DmaStartAddress = (UINT32)DeviceAddress;
    Status = PciIo->Mem.Write (
                          PciIo,
                          EfiPciIoWidthUint32,  // Width
                          1,                    // BarIndex
                          0x20,                 // Offset
                          1,                    // Count
                          &DmaStartAddress      // Buffer
                          );
    if (EFI_ERROR (Status)) {
      return Status;
    }

    //
    // Write the length of the DMA to MMIO Register 0x24 of Bar #1
    // This write operation will also start the DMA transaction
    //
    Status = PciIo->Mem.Write (
                          PciIo,
                          EfiPciIoWidthUint32,  // Width
                          1,                    // BarIndex
                          0x24,                 // Offset
                          1,                    // Count
                          &NumberOfBytes        // Buffer
                          );
    if (EFI_ERROR (Status)) {
      return Status;
    }

    //
    // Call PollMem() to poll for Bit #0 in MMIO register 0x10 of Bar #1
    //
    Status = PciIo->PollMem (
                     PciIo,
                     EfiPciIoWidthUint32,     // Width
                     1,                       // BarIndex
                     0x10,                    // Offset
                     0x01,                    // Mask
                     0x01,                    // Value
                     10000000,                // Poll for 1 second
                     &ControllerStatus        // Result
                     );
```

```
    if (EFI_ERROR (Status)) {
      return Status;
    }

    //
    // Call Unmap() to complete the bus master read operation
    //
    Status = PciIo->Unmap (PciIo, Mapping);
    if (EFI_ERROR (Status)) {
      return Status;
    }

    //
    // Update the HostAddress and Length remaining based upon the number
    // of bytes transferred
    //
    HostAddress = HostAddress + NumberOfBytes;
    *Length = *Length - NumberOfBytes;
  } while (*Length != 0);

  return Status;
)
```

Example 14-10. Bus Master Read Operation

14.5.7 DMA Bus Master Write Operation

The general algorithm for performing a bus master write operation is as follows:

- Call **Map()** with an *Operation* of **EfiPciOperationBusMasterWrite**.
- Program the DMA bus master with the *DeviceAddress* returned by **Map()**.
- Program the DMA bus master with the *NumberOfBytes* returned by **Map()**.
- Start the DMA bus master.
- Wait for the DMA bus master to complete the bus master write operation.
- Read any register on the PCI controller to flush all PCI write buffers (see the PCI specification, section 3.2.5.2). In many cases, this read is being done for other purposes, but if not, put in a dummy read.
- Call **Flush()**.
- Call **Unmap()**.
- The processor may read the contents of the DMA buffer using *HostAddress*.

The code example in Example 14-11 shows a function to perform a bus master write operation on a PCI controller. The PCI controller is accessed through the parameter *PciIo*. The system memory buffer that will be written by the bus master is specified by *HostAddress* and *Length*. This function will perform one or more bus master write operations until either *Length* bytes have been written by the bus master or an error is detected. The PCI controller in this example has three MMIO registers in BAR #1. The MMIO register at offset 0x10 is a status register that the function uses to see if the DMA operation is complete or not. The function writes the start of the DMA transaction to the MMIO register at offset 0x20 and the length of the DMA transaction to the MMIO register at offset 0x24. The write operation to offset 0x24 also starts the DMA write operation. The services of the PCI I/O Protocol that are used in this example include **Map()**, **Unmap()**, **Mem.Write()**, **PollMem()**, and **Flush()**. This example is for a 32-bit PCI bus

master. If a 64-bit PCI bus master was being used, then there would be two 32-bit MMIO registers to specify the start address and two 32-bit MMIO registers to specify the length. If the PCI bus master supports 64-bit DMA addressing, then the **EFI_PCI_ATTRIBUTE_DUAL_ADDRESS_CYCLE** attribute must be set in the **Start()** service of the PCI driver.

```
EFI_STATUS
DoBusMasterWrite (
  IN EFI_PCI_IO_PROTOCOL  *PciIo,
  IN UINT8                *HostAddress,
  IN UINTN                *Length
  )

{
  EFI_STATUS            Status;
  UINTN                 NumberOfBytes;
  EFI_PHYSICAL_ADDRESS  DeviceAddress;
  VOID                  *Mapping;
  UINT32                DmaStartAddress;
  UINT32                DummyRead;
  UINT32                ControllerStatus;

  //
  // Loop until the entire buffer specified by HostAddress and Length
  // has been written by the PCI DMA bus master
  //
  do {
    //
    // Call Map() to retrieve the DeviceAddress to use for the bus master
    // write operation. The Map() function may not support performing
    // a DMA operation for the entire length, so it may be broken up into
    // smaller DMA operations.
    //
    NumberOfBytes = *Length;
    Status = PciIo->Map (
                     PciIo,
                     EfiPciIoOperationBusMasterWrite,
                     (VOID *)HostAddress,
                     &NumberOfBytes,
                     &DeviceAddress,
                     &Mapping
                     );
    if (EFI_ERROR (Status)) {
      return Status;
    }

    //
    // Write the DMA start address to MMIO Register 0x20 of Bar #1
    //
    DmaStartAddress = (UINT32)DeviceAddress;
    Status = PciIo->Mem.Write (
                       PciIo,
                       EfiPciIoWidthUint32,  // Width
                       1,                    // BarIndex
                       0x20,                 // Offset
                       1,                    // Count
                       &DmaStartAddress      // Buffer
                       );
    if (EFI_ERROR (Status)) {
```

```
    return Status;
  }

  //
  // Write the length of the DMA to MMIO Register 0x24 of Bar #1
  // This write operation will also start the DMA transaction
  //
  Status = PciIo->Mem.Write (
                        PciIo,
                        EfiPciIoWidthUint32,  // Width
                        1,                    // BarIndex
                        0x24,                 // Offset
                        1,                    // Count
                        &NumberOfBytes        // Buffer
                        );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Call PollMem() to poll for Bit #0 in MMIO register 0x10 of Bar #1
  //
  Status = PciIo->PollMem (
                    PciIo,
                    EfiPciIoWidthUint32,     // Width
                    1,                       // BarIndex
                    0x10,                    // Offset
                    0x01,                    // Mask
                    0x01,                    // Value
                    10000000,                // Poll for 1 second
                    &ControllerStatus        // Result
                    );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Read MMIO Register 0x24 of Bar #1 to flush all posted writes
  // from the PCI bus master and through PCI-to-PCI bridges.
  //
  Status = PciIo->Mem.Read (
                        PciIo,
                        EfiPciIoWidthUint32,  // Width
                        1,                    // BarIndex
                        0x24,                 // Offset
                        1,                    // Count
                        &DummyRead            // Buffer
                        );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Call Flush() to flush all write transactions to system memory
  //
  Status = PciIo->Flush (PciIo);
  if (EFI_ERROR (Status)) {
    return Status;
  }
```

```
    //
    // Call Unmap() to complete the bus master write operation
    //
    Status = PciIo->Unmap (PciIo, Mapping);
    if (EFI_ERROR (Status)) {
      return Status;
    }

    //
    // Update the HostAddress and Length remaining based upon the number
    // of bytes transferred
    //
    HostAddress = HostAddress + NumberOfBytes;
    *Length = *Length - NumberOfBytes;
  } while (*Length != 0);

  return Status;
)
```

Example 14-11. Bus Master Write Operation

14.5.8 DMA Bus Master Common Buffer Operation

A PCI driver uses common buffers when a memory region requires simultaneous access by both the processor and a PCI bus master. A common buffer is typically allocated in the **Start()** service and freed in the **Stop()** service. This mechanism is very different from the bus master read and bus master write operations where the PCI driver transfers the ownership of a memory region from the processor to the bus master and back to the processor.

The general algorithm for allocating a common buffer in the **Start()** service is as follows:

- Call **AllocateBuffer()** to allocate a common buffer.
- Call **Map()** with an *Operation* of **EfiPciOperationBusMasterCommonBuffer**.
- Program the DMA bus master with the *DeviceAddress* returned by **Map()**.
- The common buffer can now be accessed equally by the processor (using *HostAddress*) and the DMA bus master (using *DeviceAddress*) .

The general algorithm for freeing a common buffer in the **Stop()** service is as follows:

- Call **Unmap()**.
- Call **FreeBuffer()**.

The code example in Example 14-12 shows a function that the **Start()** service may call to set up a common buffer operation on a PCI controller. The function accesses the PCI controller through the *PciIo* parameter. The function also allocates a common buffer of *Length* bytes and returns the address of the common buffer in *HostAddress*. A mapping is created for the common buffer and returned in the parameter *Mapping*. The MMIO register at offset 0x18 of BAR #1 is the start address of the common buffer from the PCI controller's perspective. The services of the PCI I/O Protocol that are used in this example include **AllocateBuffer()**, **Map()**, and **Mem.Write()**. This example is for a 32-bit PCI bus master. A 64-bit PCI bus master requires two 32-bit MMIO registers to specify the start address, and the **EFI_PCI_ATTRIBUTE_DUAL_ADDRESS_CYCLE** attribute must be set in the **Start()** service of the PCI driver.

```
EFI_STATUS
SetupCommonBuffer (
  IN EFI_PCI_IO_PROTOCOL  *PciIo,
  IN UINT8                **HostAddress,
  IN UINTN                Length,
  OUT VOID                **Mapping
  )

{
  EFI_STATUS            Status;
  UINTN                 NumberOfBytes;
  EFI_PHYSICAL_ADDRESS  DeviceAddress;
  UINT32                DmaStartAddress;

  //
  // Allocate a common buffer from anywhere in system memory of type
  // EfiBootServicesData.
  //
  Status = PciIo->AllocateBuffer (
                    PciIo,
                    AllocateAnyPages,
                    EfiBootServicesData,
                    EFI_SIZE_TO_PAGES (Length),
                    HostAddress,
                    0
                    );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Call Map() to retrieve the DeviceAddress to use for the bus master
  // common buffer operation. If the Map() function cannot support
  // a DMA operation for the entire length, then return an error.
  //
  NumberOfBytes = Length;
  Status = PciIo->Map (
                    PciIo,
                    EfiPciIoOperationBusMasterCommonBuffer,
                    (VOID *)*HostAddress,
                    &NumberOfBytes,
                    &DeviceAddress,
                    Mapping
                    );
  if (!EFI_ERROR (Status) && NumberOfBytes != Length) {
```

```
    PciIo->Unmap (PciIo, *Mapping);
    Status = EFI_OUT_OF_RESOURCES;
  }
  if (EFI_ERROR (Status)) {
    PciIo->FreeBuffer (
             PciIo,
             EFI_SIZE_TO_PAGES (Length),
             (VOID *)*HostAddress
             );
    return Status;
  }

  //
  // Write the DMA start address to MMIO Register 0x18 of Bar #1
  //
  DmaStartAddress = (UINT32)DeviceAddress;
  Status = PciIo->Mem.Write (
                       PciIo,
                       EfiPciIoWidthUint32,  // Width
                       1,                    // BarIndex
                       0x18,                 // Offset
                       1,                    // Count
                       &DmaStartAddress      // Buffer
                       );

  if (EFI_ERROR (Status)) {
    PciIo->Unmap (PciIo, *Mapping);
    PciIo->FreeBuffer (
             PciIo,
             EFI_SIZE_TO_PAGES (Length),
             (VOID *)*HostAddress
             );
  }

  return Status;
)
```

Example 14-12. Setting up a Bus Master Common Buffer Operation

The code example in Example 14-13 shows a function that the **Stop()** service may call to free a common buffer for a PCI controller. The function accesses the PCI controller through the services of the *PciIo* parameter, and the function uses these services to free the common buffer specified by *HostAddress* and *Length*. This function will undo the mapping and free the common buffer. The services of the PCI I/O Protocol that are used in this example include **Unmap()** and **FreeBuffer()**.

```
EFI_STATUS
TearDownCommonBuffer (
  IN EFI_PCI_IO_PROTOCOL  *PciIo,
  IN UINT8                *HostAddress,
  IN UINTN                Length,
  IN VOID                 *Mapping
  )

{
  EFI_STATUS            Status;

  Status = PciIo->Unmap (PciIo, Mapping);
  if (EFI_ERROR (Status)) {
    return Status;
  }

  Status = PciIo->FreeBuffer (
                    PciIo,
                    EFI_SIZE_TO_PAGES (Length),
                    (VOID *)HostAddress
                    );

  return Status;
)
```

Example 14-13. Tearing Down a Bus Master Common Buffer Operation

14.6 Device I/O Protocol

EFI drivers that follow the EFI Driver Model must not use the services of the Device I/O Protocol. Instead, the EFI driver should use the services of the PCI I/O Protocol. The Device I/O Protocol is present in a platform because a platform must be able to run EFI applications and drivers written to follow the *EFI 1.02 Specification*. EFI 1.02 drivers should be converted to EFI Driver Model drivers. This conversion means that the only consumers of the Device I/O Protocol will be EFI applications.

15
USB Driver Design Guidelines

There are several classes of USB drivers that cooperate to provide support for USB in a platform. Table 15-1 lists these USB drivers.

Table 15-1. Classes of USB Drivers

Class of Driver	Description
Host controller driver	Typically consumes the PCI I/O Protocol on the USB host controller handle and produces the USB Host Controller Protocol.
USB bus driver	Consumes the USB Host Controller Protocol and produces a child handle for each USB controller on the USB bus. Installs the Device Path Protocol and USB I/O Protocol onto each child handle.
USB device driver	Consumes the USB I/O Protocol and produces an I/O abstraction that provides services for the console devices and boot devices that are required to boot an EFI-compliant operating system.

This chapter will concentrate on how to write host controller drivers and USB device drivers. USB drivers must follow all of the general design guidelines described in chapter 5. In addition, because most USB host controllers are PCI controllers, the PCI-specific design guidelines should also be used as described in chapter 14.

Figure 15-1 below shows an example of a USB driver stack and the protocols that the USB drivers consume and produce. Because the USB hub is a special kind of device that simply acts as a signal repeater, it is not included in Figure 15-1. The protocol interfaces for the USB Host Controller Protocol, EFI USB I/O Protocol, and USB ATAPI Protocol are listed below.

Protocol Interface Structure

```
typedef struct _EFI_USB_HC_PROTOCOL {
  EFI_USB_HC_PROTOCOL_RESET             Reset;
  EFI_USB_HC_PROTOCOL_GET_STATE         GetState;
  EFI_USB_HC_PROTOCOL_SET_STATE         SetState;
  EFI_USB_HC_PROTOCOL_CONTROL_TRANSFER  ControlTransfer;
  EFI_USB_HC_PROTOCOL_BULK_TRANSFER     BulkTransfer;
  EFI_USB_HC_PROTOCOL_ASYNC_INTERRUPT_TRANSFER
                                        AsyncInterruptTransfer;
  EFI_USB_HC_PROTOCOL_SYNC_INTERRUPT_TRANSFER
                                        SyncInterruptTransfer;
  EFI_USB_HC_PROTOCOL_ISOCHRONOUS_TRANSFER
                                        IsochronousTransfer;
  EFI_USB_HC_PROTOCOL_ASYNC_ISOCHRONOUS_TRANSFER
                                        AsyncIsochronousTransfer;
  EFI_USB_HC_PROTOCOL_GET_ROOTHUB_PORT_NUMBER
                                        GetRootHubPortNumber;
```

```
  EFI_USB_HC_PROTOCOL_GET_ROOTHUB_PORT_STATUS
                                        GetRootHubPortStatus;
  EFI_USB_HC_PROTOCOL_SET_ROOTHUB_PORT_FEATURE
                                        SetRootHubPortFeature;
  EFI_USB_HC_PROTOCOL_CLEAR_ROOTHUB_PORT_FEATURE
                                        ClearRootHubPortFeature;
  UINT16                                MajorRevision;
  UINT16                                MinorRevision;
} EFI_USB_HC_PROTOCOL;


typedef struct _EFI_USB_IO_PROTOCOL {
  EFI_USB_IO_CONTROL_TRANSFER           UsbControlTransfer;
  EFI_USB_IO_BULK_TRANSFER              UsbBulkTransfer;
  EFI_USB_IO_ASYNC_INTERRUPT_TRANSFER   UsbAsyncInterruptTransfer;
  EFI_USB_IO_SYNC_INTERRPUT_TRANSFER    UsbSyncInterruptTransfer
  EFI_USB_IO_ISOCHRONOUS_TRANSFER       UsbIsochronousTransfer;
  EFI_USB_IO_ASYNC_ISOCHRONOUS_TRANSFER
                                     UsbAsyncIsochronousTransfer;
  EFI_USB_IO_GET_DEVICE_DESCRIPTOR      UsbGetDeviceDescriptor;
  EFI_USB_IO_GET_CONFIG_DESCRIPTOR      UsbGetConfigDescriptor;
  EFI_USB_IO_GET_INTERFACE_DESCRIPTOR   UsbGetInterfaceDescriptor;
  EFI_USB_IO_GET_ENDPOINT_DESCRIPTOR    UsbGetEndpointDescriptor;
  EFI_USB_IO_GET_STRING_DESCRIPTOR      UsbGetStringDescriptor;
  EFI_USB_IO_GET_SUPPORTED_LANGUAGES    UsbGetSupportedLanguages;
  EFI_USB_IO_PORT_RESET                 UsbPortReset;
} EFI_USB_IO_PROTOCOL;


typedef struct _EFI_USB_ATAPI_PROTOCOL {
  EFI_USB_ATAPI_PACKET_CMD        UsbAtapiPacketCmd;
  EFI_USB_MASS_STORAGE_RESET      UsbAtapiReset;
  UINT32                          CommandProtocol;
} EFI_USB_ATAPI_PROTOCOL;
```

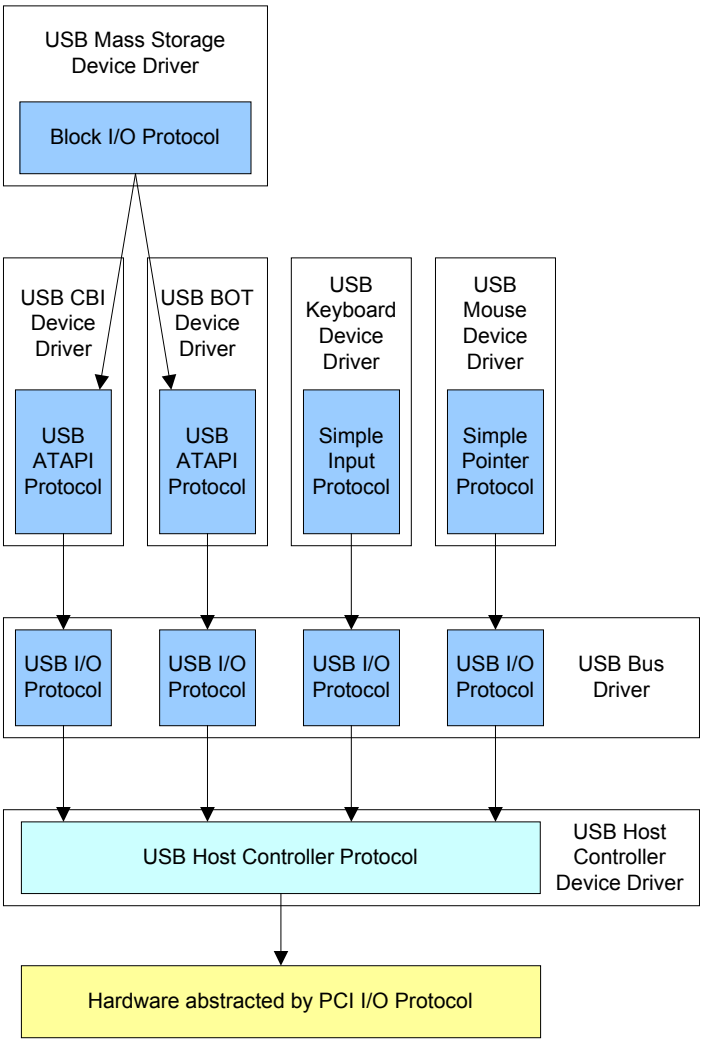

Figure 15-1. USB Driver Stack

In this example, the platform hardware produces a single USB host controller on the PCI bus. The PCI bus driver will produce a handle with **EFI_DEVICE_PATH_PROTOCOL** and **EFI_PCI_IO_PROTOCOL** installed for this USB host controller. The USB host controller driver will then consume **EFI_PCI_IO_PROTOCOL** on that USB host controller device handle and install the **EFI_USB_HC_PROTOCOL** onto the same handle.

The USB bus driver consumes the services of **EFI_USB_HC_PROTOCOL**. It uses these services to enumerate the USB bus. In this example, the USB bus driver detected a USB keyboard, a USB mouse, and two USB mass storage devices. As a result, the USB bus driver will create four *child handles* and will install the **EFI_DEVICE_PATH_PROTOCOL** and **EFI_USB_IO_PROTOCOL** onto each of those handles.

The USB mouse driver will consume the **EFI_USB_IO_PROTOCOL** and produce the **EFI_SIMPLE_POINTER_PROTOCOL**. The USB keyboard driver will consume the

EFI_USB_IO_PROTOCOL to produce the **EFI_SIMPLE_INPUT_PROTOCOL**. Because there are two types of USB command interfaces, we split the device driver in this example into two layers to support USB mass storage device. The first layer (the USB Bulk-Only Transport [BOT] driver and USB Control/Bulk/Interrupt Transport [CBI] driver) will consume the **EFI_USB_IO_PROTOCOL** and produce **EFI_USB_ATAPI_PROTOCOL**, which is a generalized interface. The second layer (the USB mass storage driver) will consume **EFI_USB_ATAPI_PROTOCOL** and produce **EFI_BLOCK_IO_PROTOCOL**.

15.1 USB Host Controller Driver

The USB host controller driver depends on which USB host controller specification that the host controller is based. Currently, the major two types of USB host controllers are the following:

- Universal Host Controller Interface (UHCI)
- Open Host Controller Interface (OHCI)

A sample driver for UHCI 1.1 is included in the *EFI 1.10 Sample Implementation.* The source code is in the directory **\EFI1.1\Edk\Drivers\Usb\Uhci**.

The USB host controller driver is a device driver, which follows the EFI Driver Model. It typically consumes the services of **EFI_PCI_IO_PROTOCOL** and produces **EFI_USB_HC_PROTOCOL**. The following sections provide guidelines for implementing the EFI Driver Binding Protocol services and USB Host Controller Protocol services for the USB host controller driver.

15.1.1 Supported()

The USB host controller driver must implement the **EFI_DRIVER_BINDING_PROTOCOL** that contains the **Supported()**, **Start()**, and **Stop()** services. The **Supported()** service evaluates the *ControllerHandle* that is passed in to check if the *ControllerHandle* represents a USB host controller that the USB host controller driver knows how to manage. The typical method of implementing this evaluation is for the USB host controller driver to retrieve the PCI configuration header from this controller and check the Class Code field and possibly other fields such as the Device ID and Vendor ID. If all these fields match the values that the USB host controller driver knows how to manage, then the **Supported()** service will return **EFI_SUCCESS**. Otherwise, the **Supported()** service will return **EFI_UNSUPPORTED**.

Example 15-1 below shows an example of the **Supported()** service for the USB host controller driver that manages a PCI controller with Class code 0x30c. First, it attempts to open the PCI I/O Protocol **BY_DRIVER**. If the PCI I/O Protocol cannot be opened, then the USB host controller driver does not support the controller that is specified by *ControllerHandle*. If the PCI I/O Protocol is opened, then the services of the PCI I/O Protocol are used to read the Class Code from the PCI configuration header. The PCI I/O Protocol is always closed with **CloseProtocol()**, and **EFI_SUCCESS** is returned if the Class Code field match.

```
//
// USB Base Class Code, Subclass Code, and Programming Interface.
//
#define PCI_CLASSC_BASE_CLASS_SERIAL              0x0c
#define PCI_CLASSC_SUBCLASS_SERIAL_USB            0x03
#define PCI_CLASSC_PI_UHCI                        0x00
```

```
//
// Class Code Register offset in PCI configuration space
//
#define CLASSC                      0x09


EFI_STATUS
UHCIDriverBindingSupported (
  IN EFI_DRIVER_BINDING_PROTOCOL                    *This,
  IN EFI_HANDLE                                     Controller,
  IN EFI_DEVICE_PATH_PROTOCOL                       *RemainingDevicePath
  )
{
  EFI_STATUS                                  Status;
  EFI_PCI_IO_PROTOCOL                         *PciIo;
  UINT8                                       UsbClassCReg[3];

  //
  // Test whether there is PCI IO Protocol attached on the
  // controller handle.
  //
  Status = gBS->OpenProtocol (
                          Controller,
                          &gEfiPciIoProtocolGuid,
                          &PciIo,
                          This->DriverBindingHandle,
                          Controller,
                          EFI_OPEN_PROTOCOL_BY_DRIVER
                          );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  Status = PciIo->Pci.Read (
PciIo,
EfiPciIoWidthUint8,
CLASSC,
3 *  sizeof(UINT8),
                            &UsbClassCReg
                            );

  if (EFI_ERROR(Status)) {

      Status = EFI_UNSUPPORTED;
      goto Exit;
  }

  //
  // Test whether the controller belongs to UHCI type
  //
  if ((UsbClassCReg[2] != PCI_CLASSC_BASE_CLASS_SERIAL)
        || (UsbClassCReg[1] != PCI_CLASSC_SUBCLASS_SERIAL_USB)
        || (UsbClassCReg.[0] != PCI_CLASSC_PI_UHCI)) {

      Status = EFI_UNSUPPORTED;
      goto Exit;
  }
```

```
Exit:
  gBS->CloseProtocol (
         Controller,
         &gEfiPciIoProtocolGuid,
         This->DriverBindingHandle,
         Controller
         );

  return Status;

}
```

Example 15-1. Supported() Service for USB Host Controller Driver

15.1.2 Start() and Stop()

The **Start()** service of the Driver Binding Protocol for the USB host controller driver also opens the PCI I/O Protocol **BY_DRIVER**. It will initialize the host controller and publish an instance of the USB Host Controller Protocol.

In today's world, some host controllers provide legacy support to be compatible with legacy devices. Under this mode, the USB input device, including mouse and keyboard, will act as if they are behind an 8042 keyboard controller. In an EFI implementation, we use native USB support instead of this legacy support. As a result, in the **Start()** service of the USB host controller driver, the USB legacy support needs to be turned off before enabling the host controller. This step is required because the legacy support will conflict with the native USB support that the EFI USB driver stack implements.

Example 15-2 shows how to turn off USB legacy support for UHCI 1.1 host controllers.

```
//
// USB legacy Support
//
#define USB_EMULATION               0xc0

VOID
TurnOffUSBLegacySupport (
  IN EFI_PCI_IO_PROTOCOL  *PciIo
  )
{
  UINT16  Command;

  //
  // Disable USB Legacy Support
  //
  Command = 0;
  PciIo->Pci.Write (
          PciIo,
          EfiPciIoWidthUint16,
          USB_EMULATION,
          1,
          &Command
         );

  return;
}
```

Example 15-2. Turning off USB Legacy Support

The **Stop()** service will do the reverse of the steps that the **Start()** service does. The USB host controller driver needs to make sure that there are no memory leaks, as well as making sure that hardware is stopped accordingly.

15.1.3 USB Host Controller Protocol Transfer Related Services

The USB Host Controller Protocol provides an I/O abstraction for a USB host controller. A USB host controller is a hardware component that interfaces to a Universal Serial Bus (USB). It moves data between system memory and devices on the Universal Serial Bus by processing data structures and generating transactions on the Universal Serial Bus. This protocol is used by a USB bus driver to perform all data transaction over the Universal Serial Bus. It also provides services to manage the USB root hub that is integrated into the USB host controller.

Example 15-3 below shows a template for the implementation of the USB Host Controller Protocol. `<<DriverName>>` denotes the name of the USB host controller driver—for example, UHCI or OHCI. `<<UsbSpecificationMajorRevision>>` denotes the major revision of the *USB Specification* that the host controller follows. For example, for *USB 1.1 Specification,* it will be 1. `<<UsbSpecificationMinorRevision>>` denotes the minor revision of that *USB Specification.* For example, for the *USB 1.1 Specification,* it will be 1.

```
EFI_USB_HC_PROTOCOL g<<DriverName>>UsbHc = {
  <<DriverName>>Reset,
  <<DriverName>>GetState,
  <<DriverName>>SetState,
  <<DriverName>>ControlTransfer,
  <<DriverName>>BulkTransfer,
  <<DriverName>>AsyncInterruptTransfer,
  <<DriverName>>SyncInterruptTransfer,
  <<DriverName>>IsochronousTransfer,
  <<DriverName>>AsyncIsochronousTransfer,
  <<DriverName>>GetRootHubPortNumber,
  <<DriverName>>GetRootHubPortStatus,
  <<DriverName>>SetRootHubPortFeature,
  <<DriverName>>ClearRootHubPortFeature,
  <<UsbSpecificationMajorRevision>>,
  <<UsbSpecificationMinorRevision>>,
};


EFI_STATUS
<<DriverName>>Reset (
  IN EFI_USB_HC_PROTOCOL            *This,
  IN UINT16                         Attributes
  )
{
}

EFI_STATUS
<<DriverName>>GetState (
  IN  EFI_USB_HC_PROTOCOL           *This,
  OUT EFI_USB_HC_STATE              *State
  )
{
```

```
}

EFI_STATUS
<<DriverName>>SetState (
  IN EFI_USB_HC_PROTOCOL            *This,
  IN EFI_USB_HC_STATE               State
  )
{
}

EFI_STATUS
<<DriverName>>ControlTransfer (
  IN     EFI_USB_HC_PROTOCOL        *This,
  IN     UINT8                      DeviceAddress,
  IN     BOOLEAN                    IsSlowDevice,
  IN     UINT8                      MaximumPacketLength,
  IN     EFI_USB_DEVICE_REQUEST     *Request,
  IN     EFI_USB_DATA_DIRECTION     TransferDirection,
  IN OUT VOID                       *Data                 OPTIONAL,
  IN OUT UINTN                      *DataLength           OPTIONAL,
  IN     UINTN                      TimeOut,
  OUT    UINT32                     *TransferResult
  )
{
}

EFI_STATUS
<<DriverName>>BulkTransfer (
  IN  EFI_USB_HC_PROTOCOL           *This,
  IN  UINT8                         DeviceAddress,
  IN  UINT8                         EndPointAddress,
  IN  UINT8                         MaximumPacketLength,
  IN OUT VOID                       *Data,
  IN OUT UINTN                      *DataLength,
  IN OUT UINT8                      *DataToggle,
  IN  UINTN                         TimeOut,
  OUT UINT32                        *TransferResult
  )
{
}

EFI_STATUS
<<DriverName>>AsyncInterruptTransfer (
  IN     EFI_USB_HC_PROTOCOL              *This,
  IN     UINT8                            DeviceAddress,
  IN     UINT8                            EndPointAddress,
  IN     BOOLEAN                          IsSlowDevice,
  IN     UINT8                            MaximumPacketLength,
  IN     BOOLEAN                          IsNewTransfer,
  IN OUT UINT8                            *DataToggle      OPTIONAL,
  IN     UINTN                            PollingInterval  OPTIONAL,
  IN     UINTN                            DataLength       OPTIONAL,
  IN     EFI_ASYNC_USB_TRANSFER_CALLBACK  CallBackFunction OPTIONAL,
  IN     VOID                             *Context         OPTIONAL
  )
{
}

EFI_STATUS
```

```
<<DriverName>>SyncInterruptTransfer (
  IN     EFI_USB_HC_PROTOCOL        *This,
  IN     UINT8                      DeviceAddress,
  IN     UINT8                      EndPointAddress,
  IN     BOOLEAN                    IsSlowDevice,
  IN     UINT8                      MaximumPacketLength,
  IN OUT VOID                       *Data,
  IN OUT UINTN                      *DataLength,
  IN OUT UINT8                      *DataToggle,
  IN     UINTN                      TimeOut,
  OUT    UINT32                     *TransferResult
  )
{
}

EFI_STATUS
<<DriverName>>IsochronousTransfer (
  IN     EFI_USB_HC_PROTOCOL        *This,
  IN     UINT8                      DeviceAddress,
  IN     UINT8                      EndPointAddress,
  IN     UINT8                      MaximumPacketLength,
  IN OUT VOID                       *Data,
  IN OUT UINTN                      DataLength,
  OUT    UINT32                     *TransferResult
  );
{
}

EFI_STATUS
<<DriverName>>AsyncIsochronousTransfer (
  IN     EFI_USB_HC_PROTOCOL          *This,
  IN     UINT8                        DeviceAddress,
  IN     UINT8                        EndPointAddress,
  IN     UINT8                        MaximumPacketLength,
  IN OUT VOID                         *Data,
  IN     UINTN                        DataLength,
  IN EFI_ASYNC_USB_TRANSFER_CALLBACK  IsochronousCallBack,
  IN VOID                             *Context   OPTIONAL
  )
{
}


EFI_STATUS
<<DriverName>>GetRootHubPortNumber (
  IN  EFI_USB_HC_PROTOCOL           *This,
  OUT UINT8                         *PortNumber
  )
{
}

EFI_STATUS
<<DriverName>>GetRootHubPortStatus (
  IN  EFI_USB_HC_PROTOCOL           *This,
  IN  UINT8                         PortNumber,
  OUT EFI_USB_PORT_STATUS           *PortStatus
  )
{
}
```

```
EFI_STATUS
<<DriverName>>SetRootHubPortFeature (
  IN EFI_USB_HC_PROTOCOL            *This,
  IN UINT8                          PortNumber,
  IN EFI_USB_PORT_FEATURE           PortFeature
  )
{
}

EFI_STATUS
<<DriverName>>ClearRootHubPortFeature (
  IN EFI_USB_HC_PROTOCOL            *This,
  IN UINT8                          PortNumber,
  IN EFI_USB_PORT_FEATURE           PortFeature
  )
{
}
```

Example 15-3. Implementing the USB Host Controller Protocol

The interfaces of USB Host Controller Protocol can be categorized into the following three aspects:

- Root hub–related services:
  - **GetRootHubPortNumber()**
  - **GetRootHubPortStatus()**
  - **SetRootHubPortFeature()**
  - **ClearRootHubPortFeature()**
- Host controller state–related services:
  - **GetState()**
  - **SetState()**
  - **Reset()**
- USB transfer–related services:
  - **ControlTransfer()**
  - **BulkTransfer()**
  - **AsyncInterruptTransfer()**
  - **SyncInterruptTransfer()**
  - **IsochronousTransfer()**
  - **AsyncIsochronousTransfer()**

For implementing root hub–related services and host controller state–related services, it mainly involves read/write operations to specific USB host controller registers. The USB host controller data sheet provides abundant information on these register usages, so this topic will not be covered in detail here.

This section concentrates on the USB transfer–related services. Those transfers can be divided into the following two categories:

- Asynchronous
- Synchronous

The asynchronous transfer means the transfer will not complete with the service's return. The synchronous transfer means that, when the service returns, the transfer will also be complete. The following sections discuss these two types of transfers in more detail.

15.1.3.1 Synchronous Transfer

The USB Host Controller Protocol provides the following four synchronous transfer services:

- **`ControlTransfer()`**
- **`BulkTransfer()`**
- **`SyncInterruptTransfer()`**
- **`IsochronousTransfer()`**

Control and bulk transfers will be done in an acceptable period of time and thus are natural synchronous transfers in the view of an EFI system. Interrupt transfers and isochronous transfers can be either asynchronous or synchronous transfers, depending on the usage model.

It is convenient for the USB drivers to use these synchronous transfer services because they do not have to worry about when the data will be ready. The transfer result will be available as soon as the function returns.

The following is an example of how to use **`BulkTransfer()`** to implement a synchronous transfer service. Generally speaking, to implement a bulk transfer service, it can be divided into the following steps:

- **Preparation:** For example, USBSTS is a status register in the USB host controller. The status register needs to be cleared before starting the control transfer.
- **Setting up the DMA direction:** By judging the end point address, the UHCI driver will decide the transfer direction and then set up the PCI bus master read or write. For example, if the transfer direction is **`EfiUsbDataIn`**, then the USB host controller will read from the DMA buffer. So the bus master write needs to be set.
- **Building the transfer context:** The *UHCI Specification* defines several structures for a transfer. For example, Queue Head (QH) and Transfer Descriptor (TD) are special structures that are used to support the requirements of control, bulk, and interrupt transfers.

  In this step, these QH and TD structures need to be created and linked to the UHCI Frame List. One possible implementation can be to create one QH and a list of TDs to form a transfer list. The QH will point to the first TD and occupy one entry in the Frame List.
- **Executing the TD and getting the result:** The USB host controller will automatically execute the TD when the timer comes. The UHCI driver needs to wait until the TDs of this transfer are all executed. After that, it will get the result of the TD execution.
- **Cleaning up:** It will delete the bulk transfer QH and TD structures from the Frame List, free related structures, and unregister the PCI DMA map.

15.1.3.2 Asynchronous Transfer

The USB Host Controller Protocol provides the following two asynchronous transfer services:

- `AsyncInterruptTransfer()`
- `AsyncIsochronousTransfer()`

To support asynchronous transfers, the USB host controller driver will register a periodical timer event. Meanwhile, it will maintain a queue for all asynchronous transfers. When the timer event is signaled, the timer event callback function will go through this queue and check whether some asynchronous transfers are complete.

Generally speaking, the main work of that timer event callback function is to go through the asynchronous transfers queue. For each asynchronous transfer, it will check whether that asynchronous transfer is completed or not and do the following:

- If it is not completed, the USB host controller driver will take no action and still keep this transfer on the queue.
- If it is completed, the USB host controller driver will copy the data that it received to a predefined data buffer and remove the related QH and TD structures. Meanwhile, it will also invoke a preregistered transfer callback function. Moreover, based on that transfer's complete status, the USB host controller driver will take different additional actions, as follows:
  - If it completed without an error, it will update the transfer data status accordingly, e.g., data toggle bit.
  - If it completed with an error, it is suggested that the USB host controller do nothing and leave the error recovery work to the related USB device driver.

15.1.3.3 Internal Memory Management

To implement USB transfers, the USB host controller driver needs to manage many small memory fragments as transfer data, for example, QH and TD. If the USB host controller driver uses the system memory management services to allocate these memory fragments each time, it is not efficient and fast enough. Thus, it is recommended that the USB host controller driver manage these kinds of internal memory usage itself. One possible implementation, as in the *EFI 1.10 Sample Implementation,* is that the host controller driver can allocate a large chunk of memory in its entry point by using EFI memory services. Then it will implement a small memory management algorithm to manage this memory to satisfy internal memory allocations. By using this simple memory management mechanism, it avoids the frequent system memory management calls.

15.1.3.4 DMA

Most USB host controllers use DMA for their data transfer between host and devices. Because the processor and USB host controller both need to access that transfer data simultaneously, the USB host controller driver shall use a common buffer for all the memory that the host controller uses for data transfer. This requirement means that the processor and the host controller have an identical view of memory. See chapter 14 for usage guidelines for the common buffer.

15.2 USB Bus Driver

The *EFI 1.10 Sample Implementation* contains a generic USB bus driver. This driver uses the services of **EFI_USB_HC_PROTOCOL** to enumerate USB devices and produce child handles with **EFI_DEVICE_PATH_PROTOCOL** and **EFI_USB_IO_PROTOCOL**.

A USB hub, including the USB root hub and common hub, is a type of USB device. The USB bus driver is responsible for the management of all USB hub devices. No device drivers are required for USB hub devices.

If EFI-based system firmware is ported to a new platform, most of the USB-related changes occur in the implementation of the USB host controller driver. Moreover, to support additional USB devices, new USB device drivers are also required. However, the USB bus driver is designed to be a generic, platform-agnostic driver. As a result, customizing the USB bus driver is strongly discouraged. The design and implementation of the USB bus driver will not be covered in detail in this document.

15.3 USB Device Driver

USB device drivers will use services provided by **EFI_USB_IO_PROTOCOL** to produce one or more protocols that provide I/O abstractions of a USB device. USB device drivers should follow the EFI Driver Model. As mentioned above, the USB device drivers will not manage hub devices because those hub devices will be managed by a USB bus driver.

15.3.1 Supported()

A USB device driver must implement the **EFI_DRIVER_BINDING_PROTOCOL** that contains the **Supported()**, **Start()**, and **Stop()** services. The **Supported()** service will check the controller handle that has been passed in to see whether this handle represents a USB device that this driver knows how to manage.

The following is the most common method for doing the check:

- Check if this handle has **EFI_USB_IO_PROTOCOL** installed. If not, this handle is not a USB device on the current USB bus.
- Get the USB interface descriptor back from this **USB_IO_DEVICE**. Check whether the values of this device's *InterfaceClass*, *InterfaceSubClass*, and *InterfaceProtocol* are identical to the corresponding values that this driver could manage.

If the above two checks are passed, it means that the USB device driver can manage the device that the controller handle represents. The **Supported()** service will return **EFI_SUCCESS**. Otherwise, the **Supported()** service will return **EFI_UNSUPPORTED**. In addition, this check process must not disturb the current state of the USB device because a different USB device driver may be controlling this USB device.

Example 15-4 is a fragment of code that shows how to implement a **Supported()** service for the USB device driver that manages a USB XYZ device with specific values for those critical fields.

```
EFI_STATUS
USBXYZDriverBindingSupported (
  IN EFI_DRIVER_BINDING_PROTOCOL    *This,
```

```
  IN EFI_HANDLE                       Controller,
  IN EFI_DEVICE_PATH_PROTOCOL         *RemainingDevicePath
  )
{
  EFI_STATUS              OpenStatus;
  EFI_USB_IO_PROTOCOL     *UsbIo;
  EFI_STATUS              Status;
  EFI_USB_INTERFACE_DESCRIPTOR    InterfaceDescriptor;

  //
  // Check if USB_IO protocol is attached on the controller handle.
  //
  OpenStatus = gBS->OpenProtocol (
                      Controller,
                      &gEfiUsbIoProtocolGuid,
                      &UsbIo,
                      This->DriverBindingHandle,
                      Controller,
                      EFI_OPEN_PROTOCOL_BY_DRIVER
                  );
  if (EFI_ERROR (OpenStatus)) {
    return OpenStatus;
  }

  //
  // Get the default interface descriptor
  //
  Status = UsbIo->UsbGetInterfaceDescriptor(
                        UsbIo,
                        &InterfaceDescriptor
                        );
  if(EFI_ERROR(Status)) {
    Status =  EFI_UNSUPPORTED;
  }    else {
    //
    // Judge whether the interface descriptor is supported by this driver
    //
    if (InterfaceDescriptor.InterfaceClass == CLASS_XYZCLASS &&
      InterfaceDescriptor.InterfaceSubClass == SUBCLASS_XYZSUBCLASS &&
      InterfaceDescriptor.InterfaceProtocol == PROTOCOL_XYZPROTOCOL) {
        Status =  EFI_SUCCESS;
    } else {
     Status =  EFI_UNSUPPORTED;
   }
 }


   gBS->CloseProtocol (
         Controller,
         &gEfiUsbIoProtocolGuid,
         This->DriverBindingHandle,
         Controller
         );

  return Status;
}
```

Example 15-4. Supported() Service of USB Device Driver

Because the **Supported()** service will be invoked many times, the USB bus driver in the *EFI 1.10 Sample Implementation* makes certain optimizations. It caches the interface descriptors, so that they do not have to read from the USB devices every time a USB device driver's **Supported()** service is invoked.

15.3.2 Start() and Stop()

The **Start()** service of the Driver Binding Protocol for a USB device driver will open the USB I/O Protocol **BY_DRIVER** and install the I/O abstraction protocol for the USB device onto the handle on which the **EFI_USB_IO_PROTOCOL** is installed.

This section provides detailed guidance on how to implement a USB device driver. It uses a USB CBI mass storage device as an example. Suppose this mass storage device has the following four endpoints:

- One control endpoint
- One interrupt endpoint
- Two bulk endpoints

For the interrupt endpoint, it is synchronous. For the bulk endpoints, one is an input endpoint and the other is an output endpoint. The following sections will cover how to implement the **Start()** and **Stop()** driver binding protocol services and USB ATAPI Protocol services.

Example 15-5 lists the private context data structure for this USB CBI device driver. The Containing Record macro (**CR()**) can be used to retrieve the private context data structure from a pointer to a produced USB ATAPI Protocol interface. See chapter 8 for more details on private context data structure design guidelines.

```
typedef struct {
  UINT32            Signature;

  EFI_USB_ATAPI_PROTOCOL          UsbAtapiProtocol;
  EFI_USB_IO_PROTOCOL             *UsbIo;

  EFI_USB_INTERFACE_DESCRIPTOR    InterfaceDescriptor;
  EFI_USB_ENDPOINT_DESCRIPTOR     BulkInEndpointDescriptor;
  EFI_USB_ENDPOINT_DESCRIPTOR     BulkOutEndpointDescriptor;
  EFI_USB_ENDPOINT_DESCRIPTOR     InterruptEndpointDescriptor;

  ... ...
} USB_CBI_DEVICE;


#define CR(Record, TYPE, Field)  \
  ((TYPE *) ( (CHAR8 *)(Record) - (CHAR8 *) &(((TYPE *) 0)->Field)))
```

Example 15-5. Implementing a USB CBI Mass Storage Device Driver

15.3.2.1 Implementing the DriverBinding.Start() Service

The code skeleton for the **Start()** service will be implemented in the following steps:

1. Open the USB I/O Protocol on *ControllerHandle* **BY_DRIVER**.
2. Get the interface descriptor using the **EFI_USB_IO_PROTOCOL.UsbGetInterfaceDescriptor()** service.
3. Prepare the USB ATAPI Protocol private data structure.

   This private data structure is in type **USB_CBI_DEVICE** and has fields for the interface descriptor, endpoint descriptor, and others.

   This step will allocate memory for this USB ATAPI Protocol private data structure and do the necessary initializations—for example, setting up the *Signature*, *UsbIo*, and *InterfaceDescriptor* fields.
4. Parse the interface descriptor.

   In this step, it will parse the *InterfaceDescriptor* that was obtained in step 2 and verify that all bulk and interrupt endpoints exit. The *NumEndpoints* field in *InterfaceDescriptor* indicates how many endpoints are in this USB interface. This code piece will first get endpoint descriptors one by one by using the **UsbGetEndpointDescriptor()** service. Then it will use the *Attributes* and *EndpointAddress* fields in *EndpointDescriptor* to judge the type of the endpoint. It will also set the endpoint descriptor to the appropriate endpoint descriptor field in the USB ATAPI Protocol private data structure.
5. Install the USB ATAPI Protocol.

15.3.2.2 DriverBinding.Stop()

The **Stop()** service will do the reverse steps as the **Start()** service. It will uninstall the USB ATAPI Protocol and close its control to the USB I/O Protocol. It will also free various allocated resources—for example, the USB ATAPI Protocol private data structure.

15.3.3 USB ATAPI Protocol Services

This section continues to use the USB mass storage device example from the last section. The USB CBI device driver will publish an instance of the USB ATAPI Protocol to abstract the mass storage device. The major service that USB ATAPI Protocol provides is to encapsulate the USB ATAPI command that the upper USB mass storage driver sends. The following will provide the code skeleton for this **UsbAtapiPacketCmd()** service.

Implement the **UsbAtapiPacketCmd()** service can be divided into the following step. This example uses a BOT device.

1. **Command phase.**

   It will send ATAPI commands through the Command Block Wrapper (CBW). The host shall send each CBW, which contains a command block, to the device via the bulk-out endpoint. The CBW shall start on a packet boundary and end as a short packet with exactly 31 (1Fh) bytes transferred. The device shall indicate a successful transport of a CBW by accepting (ACKing) the CBW.

2. **Data phase (Send/Get data)**

   It will send or get data from that USB mass storage device, based on *Direction*. If *Direction* is **EfiUsbNoData**, then no action is taken. All data transport shall begin on a packet boundary. The host shall attempt to transfer the exact number of bytes to or from the device as specified by the *dCBWDataTransferLength* and *Direction* bits.
3. **Status phase**

   It will get the status from the USB mass storage device that is packaged through the Command Status Wrapper (CSW). The device shall send each CSW to the host via the bulk-in endpoint. The CSW shall start on a packet boundary and end as a short packet with exactly 13 (Dh) bytes transferred. The CSW indicates to the host the status of the execution of the command block from the corresponding CBW. The *CSWDataResidue* field indicates how much of the data that is transferred is to be considered processed or relevant. The host shall ignore any data received beyond that which is relevant.
4. **Process status**

   It will process the status that was obtained from the previous Status phase. If any fatal error happens, it will try to recovery the USB mass storage device.

15.3.4 Asynchronous Transfer Usage

Example 15-6 shows how the USB device driver uses asynchronous transfers. It uses a USB mouse driver as an example and will use the asynchronous interrupt transfer to get mouse input.

In the USB mouse driver's Driver Binding Protocol **Start()** service, it will first initiate an asynchronous interrupt transfer.

```
Status = UsbIo->UsbAsyncInterruptTransfer(
                      UsbIo,
                      EndpointAddr,
                      TRUE,
                      PollingInterval,
                      PacketSize,
                      OnMouseInterruptComplete,
                      UsbMouseDevice
                    );
```

Example 15-6. Initiating an Asynchronous Interrupt Transfer in a USB Mouse Driver

In Example 15-7, **OnMouseInterruptComplete()** is the corresponding asynchronous interrupt transfer callback function. In this function, if the passing *Result* parameter indicates an error, it will clear the endpoint error status, unregister the previous asynchronous interrupt transfer, and initiate another asynchronous interrupt transfer. If there is no error, it will set the mouse state change indicator to **TRUE** and put the data that is read into the appropriate data structure.

```
EFI_STATUS
OnMouseInterruptComplete (
  IN  VOID        *Data,
  IN  UINTN       DataLength,
  IN  VOID        *Context,
  IN  UINT32      Result
  )
{
  USB_MOUSE_DEV       *UsbMouseDev;
  EFI_USB_IO_PROTOCOL *UsbIo;
```

```
  UINT8                EndpointAddr;
  UINT32               UsbResult;

  UsbMouseDev = (USB_MOUSE_DEV *)Context;
  UsbIo = UsbMouseDev->UsbIo;

  if (Result != EFI_USB_NOERROR) {
    if ((Result & EFI_USB_ERR_STALL) == EFI_USB_ERR_STALL) {
      EndpointAddr = UsbMouseDev->IntEndpointDescriptor->EndpointAddress;

      UsbClearEndpointHalt(
        UsbIo,
        EndpointAddr,
        &UsbResult
      );
    }

    //
    // Unregister previous asynchronous interrupt transfer
    //
    UsbIo->UsbAsyncInterruptTransfer(
             UsbIo,
             UsbMouseDev->IntEndpointDescriptor->EndpointAddress,
             FALSE,
             0,
             0,
             NULL,
             NULL
           );

    //
    // Initiate a new asynchronous interrupt transfer
    //
    UsbIo->UsbAsyncInterruptTransfer(
                      UsbIo,
                      UsbMouseDev->IntEndpointDescriptor->EndpointAddress,
                      TRUE,
                      UsbMouseDev->IntEndpointDescriptor->Interval,
                      UsbMouseDev->IntEndpointDescriptor->MaxPacketSize,
                      OnMouseInterruptComplete,
                      UsbMouseDev
                 );
    return EFI_DEVICE_ERROR;
  }


  UsbMouseDev->StateChanged = TRUE;

  //
  // Parse HID data package
  // and extract mouse movements and coordinates to UsbMouseDev
  //
  ... ...


  return EFI_SUCCESS;
}
```

Example 15-7. Completing an Asynchronous Interrupt Transfer

Example 15-8 shows the **GetMouseState()** service of the Simple Pointer In Protocol that the USB mouse driver will publish. **GetMouseState()** will not initiate any asynchronous interrupt transfer. It simply checks the mouse state change indicator. If there is mouse input, it will copy the mouse input to the passing *MouseState* data structure.

```
EFI_STATUS
GetMouseState(
  IN   EFI_SIMPLE_POINTER_PROTOCOL  *This,
  OUT  EFI_SIMPLE_POINTER_STATE     *MouseState
)
{
  USB_MOUSE_DEV           *MouseDev;


  MouseDev = USB_MOUSE_DEV_FROM_MOUSE_PROTOCOL(This);

  if (MouseDev->StateChanged == FALSE) {
    return EFI_NOT_READY;
  }

  EfiCopyMem(
    MouseState,
    &MouseDev->State,
    sizeof(EFI_SIMPLE_POINTER_STATE)
  );

  //
  // Clear previous move state
  //
  ... ...

  return EFI_SUCCESS;
}
```

Example 15-8. Retrieving Pointer Movement

15.3.5 State Machine Consideration

To implement USB device support, the USB device drivers need to maintain a state machine for their own transaction process. For example, the CBI driver for a USB mass storage device needs to maintain a tri-state machine, which contains Command->[Data]->Status states. See section 15.3.3 for a description of each state.

It should work well because it looks like a handshake process that is designed to be error free. Maintaining this state machine should provide enough and robust error handling.

However, imagine the following condition:

1. A command is sent to the device that the host needs some data from the device.
2. The device's response is too slow and it keeps NAK in its data endpoint.
3. The host sees the NAK so many times that it thinks there will be no data available from the device. It will time out this data phase operation.
4. The state machine is then in the status phase. It will ask for the status data from the device.
5. The device then sends the real data phase data to the host.

6. The host cannot understand the data from the device as status data, so it will reset the device and retry the operation.
7. The necessary components of a dead loop then exist. The final result is a hanglike system, an unusable device, or both.

How can this condition be avoided? If the device keeps NAK, it means that, sooner or later, the data will be available and no assumption can be made about the data's availability. There are some cases in which the device's response is so slow that the timeout is not enough for it to get data ready. As a result, retrying the transaction in the data phase may be necessary.

15.4 Debug Techniques

Several techniques can be used to debug the USB driver stack. The following sections describe these techniques.

15.4.1 Debug Message Output

One typical debug technique is to output the debug message. You can use the **DEBUG** macro to output debug message; see chapter 21 for the usage of the **DEBUG** macro. You can print the message both in the entry point and exit point of functions. By doing so, you can get the call stack and easily locate the error function. It is not suggested to print the debug message in a frequently called function, such as a timer handler.

15.4.2 USB Bus Analyzer

There are still some conditions where using the **DEBUG** macro is not enough for a developer to find the problem. One technique is to use a USB bus analyzer. Because a bus analyzer is inserted between the host and the device, the bus analyzer can get all the traffic on a single USB cable. With the USB bus's traffic information, some hard bugs can be root caused—for example, when a host controller loses packets on some occasions. Also, for the state machine chaos problem that was introduced in section 15.3.5, a bus analyzer will help to look at the packets' sequences and the unfinished state machine. The problem can also be quickly solved.

15.4.3 USBCheck/USBCV tool

Another useful tool for debugging is the USBCheck/USBCV tool from www.usb.org. This tool is very helpful when you want to see whether a device complies with the driver you are writing. Consider, for example, a case where a developer has written a USB imaging device driver for a generic imaging device such as a digital camera. If an end-user claims that this driver does not work for his or her specific brand of digital camera and the developer does not have such a camera on hand, the developer can ask the user to use the USBCheck/USBCV tool set and find out the device's *InterfaceClass*, *InterfaceSubClass*, and *InterfaceProtocol*. The developer can then use this information to evaluate whether the camera should be supported by the driver.

15.5 Nonconformant Device

There are always arguments on how to deal with devices that do not conform to the *USB Specification.* It is suggested to stick to the specification and reject any nonconformant devices.

However, even if the device is nonconformant and the USB driver stack should reject it, developers need to make sure that the nonconformant device will not cause any system failures. The developer cannot make any assumptions about the device's behavior. It is essential for the end-user's experience that the nonconformant device does not negatively affect the system.

16
SCSI Driver Design Guidelines

This chapter focuses on the design and implementation of EFI SCSI drivers. Most SCSI controllers are PCI controllers, and the SCSI drivers managing them are also PCI drivers. As such, they must follow all of the design guidelines described in chapter 14, as well as the general guidelines described in chapter 5. In addition, this chapter covers the guidelines that apply specifically to the management of SCSI host controllers, SCSI channels, and SCSI devices.

16.1 SCSI Driver Overview

Per the *EFI 1.10 Specification,* the EFI SCSI driver is mainly referred as a SCSI host controller driver that manages a SCSI host controller that contains one or more SCSI channels. It may create SCSI channel handles for each SCSI channel and attach the SCSI Pass Thru Protocol and Device Path Protocol to each handle that the driver produced. The SCSI driver is also responsible for the following:

- Enumerating SCSI devices that are attached on each SCSI channel
- Creating SCSI device handles for each SCSI device that is detected
- Attaching the I/O abstraction protocol instances to the child handles that represent SCSI devices

These I/O abstractions allow the SCSI device to be used in the preboot environment. These I/O abstractions, such as Block I/O, are often used to boot an EFI-compliant OS. See the *EFI 1.10 Specification* for details about **`EFI_SCSI_PASS_THRU_PROTOCOL`**.

16.2 EFI SCSI Driver on SCSI Adapters

An EFI SCSI driver follows the EFI Driver Model. Depending on the adapter that it manages, a SCSI driver can be categorized as either a bus driver or a hybrid driver. It may create child handles for each SCSI channel. A SCSI driver is typically a SCSI chip-specific driver because it has to know the details of what SCSI host adapter it is currently managing and initialize the SCSI adapter on a hardware-specific basis.

Because there may be multiple SCSI host adapters that can be managed by a single SCSI driver in the same platform, it is recommended that the SCSI host controller driver be designed to be re-entrant, as described in section 5.6 of this document.

16.2.1 EFI SCSI Driver on Single-Channel SCSI Adapter

If the target SCSI adapter supports only one channel, then the SCSI driver can simply do the following:

- Attach the **EFI_SCSI_PASS_THRU_PROTOCOL** and the **EFI_DEVICE_PATH_PROTOCOL** on the same handle as **EFI_PCI_IO_PROTOCOL**.
- Enumerate the SCSI devices on this channel.
- Create child device handles for each detected SCSI device.

The SCSI driver is also responsible for producing the Block I/O Protocol or other equivalent I/O abstraction protocols on the SCSI device handle. Figure 16-1 shows an example implementation on a single-channel SCSI adapter:

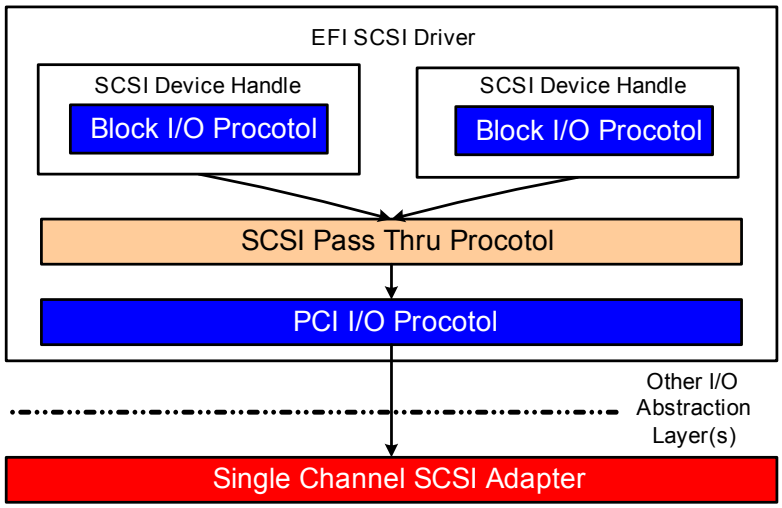

Figure 16-1. Sample SCSI Driver Implementation on Single-Channel Adapter

Because there is only one SCSI channel, the SCSI driver can simply implement one instance of the SCSI Pass Thru Protocol based on the PCI I/O Protocol that is installed by the PCI bus driver. The SCSI driver then scans the physical channel for SCSI devices. If SCSI devices are available, the SCSI driver should do the following:

- Build the corresponding child device handle and device path for each attached SCSI device.
- Implement the Block I/O Protocol instance.
- Hook it up on each child device handle.

This case is quite simple because the one SCSI pass thru–per–SCSI channel mapping is very clear.

16.2.2 EFI SCSI Driver on Multichannel SCSI Adapter

An EFI SCSI driver becomes more complex if the SCSI adapter to be managed produces multiple SCSI channels. Figure 16-2 shows a possible SCSI driver implementation on a two-channel SCSI adapter.

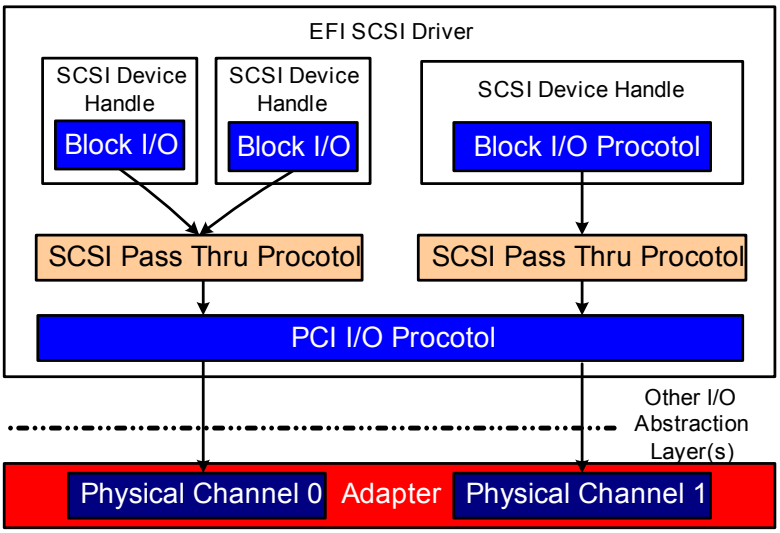

Figure 16-2. Sample SCSI Driver Implementation on a Multichannel Adapter

In this case, one SCSI adapter produces two physical SCSI channels. The SCSI driver should do the following:

- Create SCSI channel handles for each physical channel.
- Produce an instance of the SCSI Pass Thru Protocol for each of them.
- Attach all the instances of the SCSI Pass Thru Protocol to them.

The SCSI driver then enumerates each physical channel for the available SCSI devices. If there are any, the SCSI driver creates the child handle for each SCSI device and produces the block I/O for each device based on the individual SCSI Pass Thru Protocol instance.

16.2.3 EFI SCSI Driver on RAID SCSI Adapter

An EFI SCSI driver can also support SCSI adapters with RAID capability. Figure 16-3 shows an example implementation with two physical channels and one logical channel.

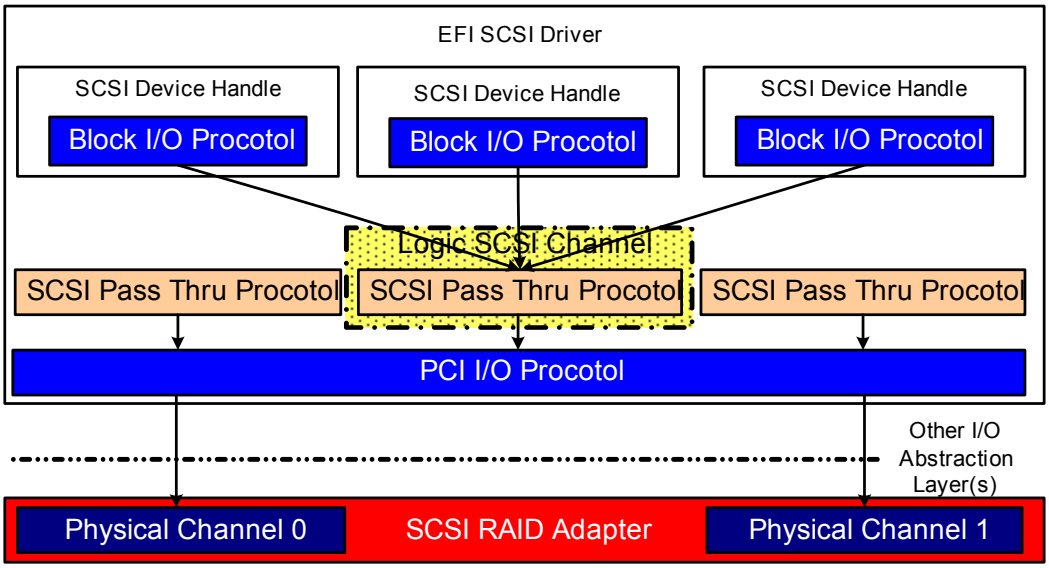

Figure 16-3. Sample SCSI Driver Implementation on Multichannel RAID Adapter

In this example, two physical channels are implemented on the SCSI adapter. The RAID component then configures these two channels to produce a logical SCSI channel. The two physical channels do have the SCSI Pass Thru installed, but the SCSI driver should not enumerate SCSI devices on physical channels or create any child handles on any physical channel. For the logical channel, the SCSI driver may produce its own SCSI Pass Thru Protocol instance based on the RAID configuration and the SCSI pass thru from physical channels or directly by the PCI I/O Protocol. SCSI devices on this logical channel should be enumerated only by an EFI SCSI driver. The EFI SCSI driver should also create child device handles and hook up the block I/O on these child handles.

The SCSI adapter hardware may not be able to expose the physical SCSI channel(s) to upper-level software when implementing RAID. If the physical SCSI channel cannot be exposed to upper software, then the SCSI driver is required only to produce a single RAID logical channel. The EFI SCSI driver will do the following:

- Scan for SCSI devices.
- Create child handles.
- Expand the device path tree based on this logical channel.

Although the basic theory is the same as the one on a physical channel, it is different from a manufacturing and diagnostic perspective. If the physical SCSI channels are exposed, the EFI SCSI driver will install the SCSI Pass Thru Protocol on each physical channel. Therefore, any SCSI command, including diagnostic ones, can be sent to an individual channel, which is very helpful on manufacturing lines. Furthermore, the diagnostic command can be sent simultaneously to all

physical channels using the nonblocking mode that is supported by SCSI Pass Thru Protocol. The diagnostic process may considerably benefit from the performance gain. In summary, it is suggested to expose physical SCSI channel whenever possible.

Of course, there are many possible designs for implementing SCSI RAID functionality. The point is that an EFI SCSI driver can be designed and implemented for a wide variety of SCSI adapters, and those EFI SCSI drivers can produce the SCSI Pass Thru Protocol for SCSI channels and generic EFI I/O abstractions such as block I/O for the SCSI devices on those SCSI channels.

16.2.4 EFI Driver Binding for EFI SCSI Driver

Like many other drivers that follow the EFI Driver Model, the image entry point of a SCSI driver installs only the Driver Binding Protocol instance on the image handle. All three of the services in the Driver Binding Protocol—**Supported()**, **Start()**, and **Stop()**—must be implemented by an EFI SCSI driver.

The **Supported()** function tests to see whether the given handle is a manageable SCSI adapter. In this function, a SCSI driver should check that **EFI_DEVICE_PATH_PROTOCOL** and **EFI_PCI_IO_PROTOCOL** are present to ensure the handle that is passed in represents a PCI device. In addition, a SCSI driver should also check the *ClassCode*, *VendorId*, and *DeviceId* that are read from the device's PCI configuration header to make sure it is a compliant SCSI adapter that can be managed by the EFI SCSI driver.

The **Start()** function tells the SCSI driver to start managing the SCSI controller. In this function, a SCSI driver should use chip-specific knowledge to do the following:

- Initialize the SCSI host controller.
- Enable the PCI device.
- Allocate resources.
- Construct data structures for the driver to use.
- Implement the interfaces that are defined in **EFI_SCSI_PASS_THRU_PROTOCOL**.

If the SCSI adapter is a single-channel adapter, then the EFI SCSI driver should install **EFI_SCSI_PASS_THRU_PROTOCOL** on the same handle that has the PCI I/O Protocol attached. If the SCSI adapter is a multichannel adapter, then the driver should also do the following:

- Enumerate the SCSI channels that are supported by the host controller.
- Create SCSI channel handles for each detected SCSI channel.
- Append the device path for each channel handle.
- Attach **EFI_DEVICE_PATH_PROTOCOL** and **EFI_SCSI_PASS_THRU_PROTOCOL** to every newly created channel handle.

Furthermore, in the **Start()** function, the SCSI driver should scan for the SCSI devices on each SCSI channel. If a request is being made to scan only one SCSI device, it should scan only for the one specified. The SCSI driver should create a device handle for the SCSI device that was found, install the **EFI_DEVICE_PATH_PROTOCOL** on each device handle, and use the services of the **EFI_SCSI_PASS_THRU_PROTOCOL** to produce additional EFI I/O abstraction protocols, such as the Block I/O Protocol. The **Start()** function should not scan for the SCSI devices every time the driver is started. It should depend on whether a full device path to a specific target is passed in or if a **NULL** device path is passed in. If a **NULL** device path is passed in, the SCSI Pass Thru

Protocol driver should create a device handle for each device that was found in the scan behind the controller. If the adapter has nonvolatile storage, the driver may also want to take advantage of the stored information on whether or not to scan the devices behind the controller.

The **Stop()** function performs the opposite operations as **Start()**. Generally speaking, a SCSI driver is required to do the following:

- Disable the SCSI adapter.
- Release all resources that were allocated for this driver.
- Close the protocol instances that were opened in the **Start()** function.
- Uninstall the protocol interfaces that were attached on the host controller handle.

In general, if it is possible to design an EFI SCSI driver to create one child at a time, it should do so to support the rapid boot capability in the EFI Driver Model. Each of the child device handles created in **Start()** must contain a Device Path Protocol instance and a SCSI I/O abstraction layer. The format of device paths for SCSI devices is described in section 16.5.

16.2.5 Implementing the SCSI Pass Thru Protocol

EFI_SCSI_PASS_THRU_PROTOCOL allows information about a SCSI channel to be collected and allows SCSI Request Packets to be sent to any SCSI devices on a SCSI channel, even if those devices are not boot devices. This protocol is attached to the device handle of each SCSI channel in a system that the protocol supports and can be used for diagnostics. It may also be used to build a block I/O driver for SCSI hard drives and SCSI CD-ROM or DVD drives to allow those devices to become boot devices.

The protocol interface for the SCSI Pass Thru Protocol is listed below:

Protocol Interface Structure

```
typedef struct _EFI_SCSI_PASS_THRU_PROTOCOL {
  EFI_SCSI_PASS_THRU_MODE                *Mode;
  EFI_SCSI_PASS_THRU_PASSTHRU            PassThru;
  EFI_SCSI_PASS_THRU_GET_NEXT_DEVICE     GetNextDevice;
  EFI_SCSI_PASS_THRU_BUILD_DEVICE_PATH   BuildDevicePath;
  EFI_SCSI_PASS_THRU_GET_TARGET_LUN      GetTargetLun;
  EFI_SCSI_PASS_THRU_RESET_CHANNEL       ResetChannel;
  EFI_SCSI_PASS_THRU_RESET_TARGET        ResetTarget;
} EFI_SCSI_PASS_THRU_PROTOCOL;
```

For a detailed description of **EFI_SCSI_PASS_THRU_PROTOCOL**, see chapter 13 of the *EFI 1.10 Specification.*

Before implementing the SCSI Pass Thru Protocol, the SCSI driver should configure the SCSI core to a defined state. In practice, the SCSI adapter usually maps a set of SCSI core registers in I/O or memory-mapped I/O space. Although the detailed layout or functions of these registers vary from one SCSI hardware to another, the SCSI driver should use specific knowledge to set up the proper SCSI working mode (SCSI-I, SCSI-II, Ultra SCSI, and so on) and configure the timing registers for the current mode. Other considerations include parity options, DMA engine and interrupt initialization, among others.

All the hardware-related settings should be completed before any SCSI Pass Thru Protocol function is called. The initialization is better accomplished in the Driver Binding Protocol's **Start()** function of the SCSI controller driver, prior to hooking up the SCSI Pass Thru Protocol functions.

EFI_SCSI_PASS_THRU_PROTOCOL.*Mode* is a structure that describes the intrinsic attributes of the SCSI Pass Thru Protocol instance. Note that a non-RAID SCSI channel should set both the physical and logical attributes. A physical channel on the RAID adapter should set only the physical attribute, and the logical channel on the RAID adapter should set only the logical attribute. If the channel supports nonblocking I/O, the nonblocking attribute should also be set. Example 16-1 shows how to set those attributes on a non-RAID SCSI adapter that supports nonblocking I/O.

```
#define SCSI_CHANNEL_NAME     "Sample Channel"
#define SCSI_CONTROLLER_NAME  "Sample SCSI Adapter"
// ... .. ...
ScsiPassThruMode.ControllerName                 = SCSI_CONTROLLER_NAME;
ScsiPassThruMode.ChannelName                    = SCSI_CHANNEL_NAME;
AtapiScsiPrivate->ScsiPassThruMode.AdapterId = 4; // Target Channel Id
ScsiPassThruMode.Attributes = EFI_SCSI_PASS_THRU_ATTRIBUTES_PHYSICAL
| EFI_SCSI_PASS_THRU_ATTRIBUTES_LOGICAL
| EFI_SCSI_PASS_THRU_ATTRIBUTES_NONBLOCKIO;
ScsiPassThruMode.IoAlign     = 0; // Do not have any alignment requirement
```

Example 16-1. SCSI Pass Thru Mode Structure on Single-Channel SCSI Adapter

Example 16-2 shows how to set the SCSI *Mode* structure on a multichannel non-RAID adapter. The example fits for either channel in Figure 16-2.

```
#define SCSI_CHANNEL_NAME     "Sample Phy Channel"
#define SCSI_CONTROLLER_NAME  "Sample Multichannel SCSI Adapter"
// ... .. ...
ScsiPassThruMode.ControllerName                 = SCSI_CONTROLLER_NAME;
ScsiPassThruMode.ChannelName                    = SCSI_CHANNEL_NAME;
AtapiScsiPrivate->ScsiPassThruMode.AdapterId = 2;
// The channel does not support nonblocking I/O
ScsiPassThruMode.Attributes = EFI_SCSI_PASS_THRU_ATTRIBUTES_PHYSICAL
| EFI_SCSI_PASS_THRU_ATTRIBUTES_LOGICAL;
ScsiPassThruMode.IoAlign     = 2; // Data must be alligned on 4-byte boundary
```

Example 16-2. SCSI Pass Thru Mode Structure on Multichannel SCSI Adapter

For the RAID adapter shown in Figure 16-3, the corresponding *Mode* structures for both the physical and logical channel may be filled as shown in Example 16-3.

```
#define SCSI_PHY_CHANNEL_NAME     "Sample Phy Channel"
#define SCSI_LGC_CHANNEL_NAME     "Sample Lgc Channel"
#define SCSI_PHY_CONTROLLER_NAME  "Sample RAID SCSI Adapter"
// ... .. ... Physical Channel ... .. ...
ScsiPassThruMode.ControllerName                 = SCSI_CONTROLLER_NAME;
ScsiPassThruMode.ChannelName                    = SCSI_PHY_CHANNEL_NAME;
AtapiScsiPrivate->ScsiPassThruMode.AdapterId = 0;
ScsiPassThruMode.Attributes = EFI_SCSI_PASS_THRU_ATTRIBUTES_PHYSICAL
| EFI_SCSI_PASS_THRU_ATTRIBUTES_NONBLOCKIO;
ScsiPassThruMode.IoAlign     = 0;

//  ... .. ... Logical Channel ... .. ...
ScsiPassThruMode.ControllerName                 = SCSI_CONTROLLER_NAME;
ScsiPassThruMode.ChannelName                    = SCSI_LGC_CHANNEL_NAME;
AtapiScsiPrivate->ScsiPassThruMode.AdapterId = 2;
```

```
ScsiPassThruMode.Attributes = EFI_SCSI_PASS_THRU_ATTRIBUTES_LOGICAL
| EFI_SCSI_PASS_THRU_ATTRIBUTES_NONBLOCKIO;
ScsiPassThruMode.IoAlign     = 0;
```

Example 16-3. SCSI Pass Thru Mode Structures on RAID SCSI Adapter

The **EFI_SCSI_PASS_THRU_PROTOCOL.GetNextDevice()** and **EFI_SCSI_PASS_THRU_PROTOCOL.GetTargetLun()** functions provide a way to walk on different devices within a channel. The SCSI controller driver may implement it by internally maintaining an active device flag. Use this flag and channel-specific knowledge to figure out what device is next, as well as what device is first.

The **EFI_SCSI_PASS_THRU_PROTOCOL.BuildDevicePath()** function facilitates the construction of a SCSI device path. The device path for the SCSI device can be many kinds because of enormous device types that are supported by the SCSI pass thru mechanism. The detailed device category can be identified only by the SCSI pass thru implementation, which is why the function goes into the SCSI Pass Thru Protocol. See the last section in this chapter for device path examples for SCSI devices.

The **EFI_SCSI_PASS_THRU_PROTOCOL.PassThru()** function is the most important function when implementing the SCSI Pass Thru Protocol. In this function, the SCSI pass thru driver should do the following:

- Initialize the internal register for command/data transfer.
- Put valid SCSI packets into hardware-specific memory or register locations.
- Start the transfer.
- Optionally wait for completion of the execution.

The better error handling mechanism in this function helps to develop a more robust driver. Although most SCSI adapters support both synchronous and asynchronous data transfers, some may not support synchronous mode. In this case, the SCSI driver may implement the blocking SCSI I/O that is required by the *EFI 1.10 Specification* using the polling mechanism. Polling can be based on a timer interrupt or simply by polling the internal register. Do not return until all I/O requests are completed or else an unhandled error is encountered.

Example 16-4 below shows a template for implementing the SCSI Pass Thru Protocol.

```
// In the SCSI Channel private context data definition
typedef struct {
    UINTN                        Signature;
    EFI_HANDLE                   Handle;
    EFI_SCSI_PASS_THRU_PROTOCOL  ScsiPassThru;
    EFI_SCSI_PASS_THRU_MODE      ScsiPassThruMode;
    // .. .. ..
} <<DriverName>>_SCSI_PASS_THRU_CONTEXT_DATA;

// In the SCSI driver implementation file
<<DriverName>>_SCSI_PASS_THRU_CONTEXT_DATA   *ScsiContextData;

Status = gBS->AllocatePool (
                  EfiBootServicesData,
                  sizeof(<<DriverName>>_SCSI_PASS_THRU_CONTEXT_DATA),
                 (VOID **)&ScsiContextData
                 );
if (EFI_ERROR(Status)) {
```

```
   return Status;
}
EfiZeroMem (
   ScsiContextData,
   sizeof (<<DriverName>>_SCSI_PASS_THRU_CONTEXT_DATA)
   );
ScsiContextData->Signature = <<DriverName>>_SCSI_PASS_THRU_DEV_SIGNATURE;
//
// Initialize SCSI Pass Thru Protocol interface
//
ScsiContextData->ScsiPassThru.Mode = &ScsiContextData->ScsiPassThruMode;
ScsiContextData->ScsiPassThru.PassThru = <<DriverName>>ScsiPassThruPassThru;
ScsiContextData->ScsiPassThru.GetNextDevice =
               <<DriverName>>ScsiPassThruGetNextDevice;
ScsiContextData->ScsiPassThru.BuildDevicePath  = <<DriverName>>ScsiPassThru
               BuildDevicePath;
AtapiScsiPrivate->ScsiPassThru.GetTargetLun     =
               <<DriverName>>ScsiPassThruGetTargetLun;
ScsiContextData->ScsiPassThru.ResetChannel      =
               <<DriverName>>ScsiPassThruResetChannel;
ScsiContextData->ScsiPassThru.ResetTarget       =
               <<DriverName>>ScsiPassThruResetTarget;

// In the SCSI Pass Thru implementation file
EFI_STATUS
<<DriverName>>ScsiPassThruPassThru (
  IN EFI_SCSI_PASS_THRU_PROTOCOL                  *This,
  IN UINT32                                       Target,
  IN UINT64                                       Lun,
  IN OUT EFI_SCSI_PASS_THRU_SCSI_REQUEST_PACKET   *Packet,
  IN EFI_EVENT                                    Event OPTIONAL
  )
{
}
EFI_STATUS
<<DriverName>>ScsiPassThruGetNextDevice (
  IN  EFI_SCSI_PASS_THRU_PROTOCOL                 *This,
  IN OUT UINT32                                   *Target,
  IN OUT UINT64                                   *Lun
  )
{
}
EFI_STATUS
<<DriverName>>ScsiPassThruBuildDevicePath (
  IN     EFI_SCSI_PASS_THRU_PROTOCOL              *This,
  IN     UINT32                                   Target,
  IN     UINT64                                   Lun,
  IN OUT EFI_DEVICE_PATH_PROTOCOL                 **DevicePath
  )
{
}
EFI_STATUS
<<DriverName>>ScsiPassThruGetTargetLun (
  IN  EFI_SCSI_PASS_THRU_PROTOCOL                 *This,
  IN  EFI_DEVICE_PATH_PROTOCOL                    *DevicePath,
  OUT UINT32                                      *Target,
  OUT UINT64                                      *Lun
  )
{
```

```
}
EFI_STATUS
<<DriverName>>ScsiPassThruResetChannel (
  IN  EFI_SCSI_PASS_THRU_PROTOCOL                 *This
  )
{
}
EFI_STATUS
<<DriverName>>ScsiPassThru ResetTarget (
  IN EFI_SCSI_PASS_THRU_PROTOCOL                  *This,
  IN UINT32                                       Target,
  IN UINT64                                       Lun
  )
{
}
```

Example 16-4. SCSI Pass Thru Protocol Template

16.3 Implementing SCSI Pass Thru Protocol on a SCSI Command Set–Compatible Device

The SCSI Pass Thru Protocol defines a method to directly access SCSI devices. This protocol provides interfaces that allow a generic driver to produce the Block I/O Protocol for SCSI disk devices and allows an EFI utility to issue commands to any SCSI device. The main reason to provide such an access is to enable S.M.A.R.T. functionality during POST (i.e., issuing Mode Sense, Mode Select, and Log Sense to SCSI devices). This enabling is accomplished using the generic interfaces that are defined in the SCSI Pass Thru Protocol. The implementation of this protocol will also enable additional functionality in the future without modifying the SCSI drivers that are built on top of the SCSI driver. Furthermore, the SCSI Pass Thru Protocol is not limited to SCSI adapters. It is applicable to all channel technologies that use SCSI commands such as ATAPI, iSCSI, and Fibre Channel. This section shows some examples that demonstrate how to implement the SCSI Pass Thru Protocol on SCSI command set–compatible technology.

16.3.1 SCSI Pass Thru Protocol on ATAPI

This section shows how to implement the SCSI Pass Thru Protocol for ATAPI devices.

Decode the **(Target, Lun)** pair using the intrinsic property of the technology or device. In this example, ATAPI supports only four devices, so the **(Target, Lun)** pair can be decoded by determining the IDE channel (primary/secondary) and IDE device (master/slave).

If the corresponding technology or device supports the channel reset operation, use it to implement **EFI_SCSI_PASS_THRU_PROTOCOL.ResetChannel()**; if not, it may be implemented by resetting all attached devices on the channel and re-enumerating them.

In the **EFI_SCSI_PASS_THRU_PROTOCOL.BuildDevicePath()** function, all target devices should be built on a node based on the channel knowledge. Example 16-5 shows an ATAPI node being built.

```
typedef struct {
    UINTN                           Signature;
    EFI_HANDLE                      Handle;
```

```
    EFI_SCSI_PASS_THRU_PROTOCOL   ScsiPassThru;
    EFI_SCSI_PASS_THRU_MODE       ScsiPassThruMode;
    EFI_PCI_IO_PROTOCOL           *PciIo;
    IDE_BASE_REGISTERS            *IoPort;
    CHAR16                        ControllerName[100];
    CHAR16                        ChannelName[100];
    UINT32                        LatestTargetId;
    UINT64                        LatestLun;
} ATAPI_SCSI_PASS_THRU_DEV;

EFI_STATUS
SampleAtapiScsiPassThruBuildDevicePath (
  IN     EFI_SCSI_PASS_THRU_PROTOCOL    *This,
  IN     UINT32                         Target,
  IN     UINT64                         Lun,
  IN OUT EFI_DEVICE_PATH_PROTOCOL       **DevicePath
  )
{
  ATAPI_SCSI_PASS_THRU_DEV    *AtapiScsiPrivate;
  EFI_DEV_PATH                *Node;
  EFI_STATUS                  Status;

  // Checking parameters..........

  // Retrieve Device Private Data Structure via _CR() macro
  AtapiScsiPrivate = ATAPI_SCSI_PASS_THRU_DEV_FROM_THIS (This);

  Status = gBS->AllocatePool (
                    EfiBootServicesData,
                    sizeof(EFI_DEV_PATH),
                    (VOID **)&Node);
  if (EFI_ERROR(Status)) {
    return EFI_OUT_OF_RESOURCES;
  }

  gBS->SetMem (Node, sizeof(EFI_DEV_PATH), 0);
  Node->DevPath.Type           = MESSAGING_DEVICE_PATH;
  Node->DevPath.SubType        = MSG_ATAPI_DP;
  (&Node->DevPath)->Length[0]  = (UINT8) (sizeof(ATAPI_DEVICE_PATH));
  (&Node->DevPath)->Length[1]  = (UINT8) ((sizeof(ATAPI_DEVICE_PATH)) >> 8);
  Node->Atapi.PrimarySecondary = (UINT8)(Target / 2);
  Node->Atapi.SlaveMaster      = (UINT8)(Target % 2);
  Node->Atapi.Lun              = (UINT16)Lun;

  *DevicePath = (EFI_DEVICE_PATH_PROTOCOL*)Node;
  return EFI_SUCCESS;
}
```

Example 16-5. Building Device Path for ATAPI Device

For the most important function, **`EFI_SCSI_PASS_THRU_PROTOCOL.PassThru()`**, it should be implemented by technology-dependent means. In this example, ATAPI supports a SCSI command using the IDE "Packet" command. Because the IDE command is delivered through a group of I/O registers, the main body of the implementation is filling the SCSI command structure to these I/O registers and then waiting for the command completion. A complete code example for the blocking I/O **`EFI_SCSI_PASS_THRU_PROTOCOL`** function can be found in the **`\EFI1.1\Edk\Drivers\AtapiPassThru`** directory of the *EFI Sample Implementation.*

For the nonblocking I/O **EFI_SCSI_PASS_THRU_PROTOCOL** function, the SCSI driver should simply submit the SCSI command and return. It may choose to poll an internal timer event to check whether the submitted command completes its execution. If so, it should signal the client event. The EFI firmware will then schedule to invoke the notification function of the client event.
Example 16-6 shows a sample nonblocking SCSI Pass Thru Protocol implementation.

```
EFI_STATUS
SampleNonBlockingScsiPassThruFunction (
  IN EFI_SCSI_PASS_THRU_PROTOCOL          *This,
  IN UINT32                               Target,
  IN UINT64                               Lun,
  IN OUT EFI_SCSI_PASS_THRU_SCSI_REQUEST_PACKET   *Packet,
  IN EFI_EVENT                            Event OPTIONAL
  )
{
  ATAPI_SCSI_PASS_THRU_DEV        *AtapiScsiPrivate;
  EFI_EVENT                       InternalEvent;
  EFI_STATUS                      Status;

  AtapiScsiPrivate = ATAPI_SCSI_PASS_THRU_DEV_FROM_THIS (This);

  //
  // Do parameter checking required by EFI specification
  //
  //..................................

  //
  // Create internal timer event in order to poll the completion. The event
  // can also be created outside of this function to avoid frequent event
  // construction/destruction.
  //
  Status = gBS->CreateEvent (
                 EFI_EVENT_TIMER | EFI_EVENT_NOTIFY_SIGNAL,
                 EFI_TPL_CALLBACK,
                 ScsiPassThruPollEventNotify,
                 AtapiScsiPrivate,
                 &InternalEvent
                 );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Signal the polling event every 200 ms. Select the interval according
  // to the specific requirement and technology.
  //
  Status = gBS->SetTimer (InternalEvent, TimerPeriodic, 2000000);
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Just submit SCSI I/O command through IDE I/O registers and return
  //
  Status = SubmitBlockingIoCommand (AtapiScsiPrivate, Target, Packet);
  return Status;
}

VOID
```

```
ScsiPassThruPollEventNotify (
  IN  EFI_EVENT                Event,
  IN  VOID                     *Context
  )
{
  ATAPI_SCSI_PASS_THRU_DEV     *AtapiScsiPrivate;
  BOOLEAN                      CommandCompleted;

  ASSERT (Context);
  AtapiScsiPrivate = (ATAPI_SCSI_PASS_THRU_DEV *)Context;
  CommandCompleted = FALSE;

  //
  // Use specific knowledge to identify whether command execution completes
  // or not. If so, set CommandCompleted as TRUE.
  //
  // ......................

  if (CommandCompleted) {
    //
    // Get client event handle from private context data structure.
    // Signal it.
    //
    gBS->SignalEvent (ClientEvent);
  }
}
```

Example 16-6. Sample Nonblocking SCSI Pass Thru Protocol Implementation

16.4 General Considerations for Developing EFI SCSI Drivers

16.4.1 SCSI Channel Enumeration

The purpose of the SCSI channel enumeration is to scan for the SCSI devices that are attached to a specific SCSI channel. Although the detailed SCSI device discovery algorithm may vary from one implementation to another, the enumeration framework is generic, not only for native SCSI adapters, but also for any SCSI-command-compatible technology such as ATAPI, iSCSI, and Fibre Channel. Example 16-7 shows a possible channel enumeration framework.

```
EFI_STATUS
ScsiChannelEnumeration (
  IN EFI_DRIVER_BINDING_PROTOCOL  *This,
  IN EFI_HANDLE                   ControllerHandle,
  IN EFI_DEVICE_PATH_PROTOCOL     *RemainingDevicePath,
  IN EFI_SCSI_PASS_THRU_PROTOCOL  *ScsiPassThru,
  IN EFI_DEVICE_PATH_PROTOCOL     *ParentDevicePath
  )
{
  EFI_STATUS                  Status;
  UINT32                      ChannelAttribute;
  UINT32                      StartPun = 0;
  UINT64                      StartLun = 0;
  UINT32                      Pun;
  UINT64                      Lun;
  BOOLEAN                     ScanOtherPuns;
```

```
  ChannelAttribute = ScsiPassThru->Mode->Attributes;
  if (((ChannelAttribute & EFI_SCSI_PASS_THRU_ATTRIBUTES_PHYSICAL) != 0) &&
      ((ChannelAttribute & EFI_SCSI_PASS_THRU_ATTRIBUTES_LOGICAL)  == 0) ) {
    //
    // This channel has just PHYSICAL attribute set. Do not do enumeration
    // on this kind of channel.
    //
    return EFI_UNSUPPORTED;
  }

  if (RemainingDevicePath == NULL) {
    StartPun = 0xFFFFFFFF;
    StartLun = 0;
  } else {
    ScsiPassThru->GetTargetLun (
                        ScsiPassThru,
                        RemainingDevicePath,
                        &StartPun,
                        &StartLun
                        );
  }

  for (Pun = StartPun, ScanOtherPuns = TRUE; ScanOtherPuns == TRUE;) {
    if (StartPun == 0xFFFFFFFF) {
      //
      // RemainingDevicePath is NULL. Scan all the possible Puns in the
      // channel.
      //
      Status = ScsiPassThru->GetNextDevice (ScsiPassThru, &Pun, &Lun);
      if (EFI_ERROR(Status)) {
        //
        // No legal Pun and Lun found anymore
        //
        break;
      }
    } else {
      //
      // Remaining Device Path is not NULL. Only scan the specified Pun.
      //
      Pun = StartPun;
      Lun = StartLun;
      ScanOtherPuns = FALSE;
    }

    //
    // Avoid creating handle for the host adapter.
    //
    if (Pun == ScsiPassThru->Mode->AdapterId) {
      continue;
    }

    //
    // Use specific knowledge discover the presence of SCSI device on
    // target (Pun, Lun). Some SCSI commands, such as INQUERY,
    // TEST_UNIT_READY, REQUEST_SENSE can be used here to detect the
    // SCSI device presence.
    //
    Status = DiscoverScsiDevice (
                 This,
```

```
                  ControllerHandle,
                  ScsiPassThru,
                  Pun,
                  Lun
                  );
    if (EFI_ERROR(Status)) {
      //
      // SCSI device doesn't exist on target (Pun, Lun);
      // continue on next one
      //
      continue;
    }

    //
    // SCSI device exists on target (Pun, Lun). Create and set up child
    // device
    //
    Status = CreateChildScsiDevice (
                  This,
                  ControllerHandle,
                  ScsiPassThru,
                  Pun,
                  Lun,
                  ParentDevicePath
                  );
    if (EFI_ERROR(Status)) {
      return Status;
    }
  }

  return Status;
}
```

Example 16-7. SCSI Channel Enumeration

16.4.2 Create SCSI Device Child Handle

It is recommended that the SCSI driver constructs a private context structure for each enumerated SCSI device. See chapter 8 in this document for the advantage of using such a private context structure. Specifically, the SCSI driver should store all required information for the child SCSI device in this data structure, including the signature, child handle value, position indicator (Pun, Lun), device type, device version, and its device path. This private context structure can be accessed via the Record macro **_CR()**, which can also be found in chapter 8 of this document.

It is also the SCSI driver's responsibility to do the following:

- Build the appropriate device path for the enumerated SCSI device.
- Install the generic EFI I/O abstraction protocol and Device Path Protocol on the newly created child SCSI device handle.

Example 16-8 shows the child handle creation process for a detected SCSI device.

```
typedef struct {
  UINT32                        Signature;
  EFI_HANDLE                    Handle;
  EFI_SCSI_IO_PROTOCOL          ScsiIo;
  EFI_DEVICE_PATH_PROTOCOL      *DevicePath;
  EFI_SCSI_PASS_THRU_PROTOCOL   *ScsiPassThru;
```

```
  UINT32                          Pun;
  UINT64                          Lun;
  UINT8                           ScsiDeviceType;
  UINT8                           ScsiVersion;
  BOOLEAN                         RemovableDevice;
} SCSI_DEVICE_CONTEXT;

EFI_STATUS
CreateChildScsiDevice (
  EFI_DRIVER_BINDING_PROTOCOL   *This,
  EFI_HANDLE                    ControllerHandle,
  EFI_SCSI_PASS_THRU_PROTOCOL   *ScsiPassThru,
  UINT32                        Pun,
  UINT64                        Lun,
  EFI_DEVICE_PATH_PROTOCOL      *ParentDevicePath
  )
{
  EFI_STATUS                    Status;
  SCSI_DEVICE_CONTEXT           *ScsiDevice;
  EFI_DEVICE_PATH_PROTOCOL      *ScsiDevicePath;

  //
  // Construct and initialize SCSI Device Private Context Structure
  //
  Status = gBS->AllocatePool (
                    EfiBootServicesData,
                    sizeof(SCSI_DEVICE_CONTEXT),
                   (VOID **)&ScsiDevice);
  if (EFI_ERROR(Status)) {
    return Status;
  }


  EfiZeroMem (ScsiDevice, sizeof(SCSI_DEVICE_CONTEXT));

  ScsiDevice->Signature    = SCSI_DEVICE_SIGNATURE;
  ScsiDevice->ScsiPassThru = ScsiPassThru;
  ScsiDevice->Pun          = Pun;
  ScsiDevice->Lun          = Lun;

  //
  // Set Device Path
  //
  Status = ScsiDevice->ScsiPassThru->BuildDevicePath (
              ScsiDevice->ScsiPassThru,
              ScsiDevice->Pun,
              ScsiDevice->Lun,
              &ScsiDevicePath
              );
  if (EFI_ERROR(Status)) {
    gBS->FreePool (ScsiDevice);
    return Status;
  }

  ScsiDevice->DevicePath = EfiAppendDevicePathNode (
                                  ParentDevicePath,
                                  ScsiDevicePath
                                  );
```

```
  //
  // The memory for ScsiDevicePath is allocated in
  // ScsiPassThru->BuildDevicePath()function and no longer used
  // after EfiAppendDevicePathNode, so free it
  //
  gBS->FreePool (ScsiDevicePath);

  if (ScsiDevice->DevicePath == NULL) {
    gBS->FreePool (ScsiDevice);
    return EFI_OUT_OF_RESOURCES;
  }

  Status = gBS->InstallMultipleProtocolInterfaces (
                    &ScsiDevice->Handle,
                    &gEfiDevicePathProtocolGuid,
                    ScsiDevice->DevicePath,
                    NULL
                 );

  if (EFI_ERROR(Status)) {
    gBS->FreePool (ScsiDevice);
  } else {
    gBS->OpenProtocol (
            ControllerHandle,
            &gEfiScsiPassThruProtocolGuid,
            (VOID **)&ScsiPassThru,
            This->DriverBindingHandle,
            ScsiDevice->Handle,
            EFI_OPEN_PROTOCOL_BY_CHILD_CONTROLLER
            );
  }
  return EFI_SUCCESS;
}
```

Example 16-8. Sample Creation of SCSI Device Child Handle

16.4.3 Produce Block I/O Protocol

One important feature for a SCSI driver is to wrap SCSI I/O capability into other standard EFI I/O abstractions that can be used by upper drivers, such as file system drivers, or by an EFI-aware OS loader. This section gives an example to demonstrate how to implement **EFI_BLOCK_IO_PROTOCOL** through the SCSI Pass Thru Protocol abstraction layer.

Because many SCSI devices are hot pluggable, one major concern when implementing Block I/O Protocol functions is that they should determine whether the represented SCSI device is available before performing any actual I/O operation. Furthermore, even if the SCSI device is available, the media in it may have been changed to another one. This scenario should also be taken into account. A SCSI device driver should use the SCSI **INQUERY** command to get the device type and version information after detecting that a SCSI device is present. If it is a removable device, send the SCSI commands **TEST_UNIT_READY** and **REQUEST_SENSE** to determine the device state and media status. If the device or media is changed, call the EFI Boot Service **ReinstallProtocolInterface()** to make the internal handle/protocol database consistent with the new device or media. Example 16-9 demonstrates this process.

```
EFI_STATUS
ScsiDiskReadBlocks (
  IN  EFI_BLOCK_IO_PROTOCOL   *This,
  IN  UINT32                  MediaId,
  IN  EFI_LBA                 LBA,
  IN  UINTN                   BufferSize,
  OUT VOID                    *Buffer
  )
{
  SCSI_XYZ_DEV_PRIVATE_DATA  *ScsiDevice;

  // ....................................

  ScsiDevice = CR(This, SCSI_XYZ_DEV_PRIVATE_DATA,
BlkIo,SCSI_XYZ_DEV_SIGNATURE)

  //
  // Use INQUERY command to get device type. This info can also be stored in
  // SCSI device private context data structure
  //
  if (!IsDeviceFixed (ScsiDevice)) {

    //
    // Detect media change by sending TEST_UNIT_READY and REQUEST_SENSE
    //
    Status = ScsiDetectMedia (ScsiDevice, FALSE, &MediaChange);
    if (EFI_ERROR(Status)) {
      return EFI_DEVICE_ERROR;
    }
    if (MediaChange) {
      gBS->ReinstallProtocolInterface (
             ScsiDevice->Handle,
             &gEfiBlockIoProtocolGuid,
             &ScsiDevice->BlkIo,
             &ScsiDevice->BlkIo
             );
    }
  }

  //
  // Perform actual I/O operation via SCSI Pass Thru Protocol here
  // ....................................
  //
  return Status;
}
```

Example 16-9. Handle Device and Media Change

Another consideration when implementing Block I/O Protocol functions is to handle an unstable storage device. You can never assume that accessing media is perfectly stable or smooth. A better error handling mechanism results in a more robust SCSI driver.

16.5 SCSI Device Path

The SCSI driver described in this document can support a SCSI channel that is generated or emulated by multiple architectures, such as SCSI-I, SCSI-II, SCSI-III, ATAPI, Fibre Channel,

iSCSI, and other future channel types. This section describes four example device paths, including SCSI, ATAPI, and Fibre Channel device paths.

16.5.1 SCSI Device Path Example

Table 16-1 shows an example device path for a SCSI device controller on a desktop platform. This SCSI device controller is connected to a SCSI channel that is generated by a PCI SCSI host controller. The PCI SCSI host controller generates a single SCSI channel, is located at PCI device number 0x07 and PCI function 0x00, and is directly attached to a PCI root bridge. The SCSI device controller is assigned SCSI ID 2, and its LUN is 0.

This sample device path consists of an ACPI device path node, a PCI device path node, a SCSI node, and a device path end structure. The _HID and _UID must match the ACPI table description of the PCI root bridge. The following is the shorthand notation for this device path:

```
ACPI(PNP0A03,0)/PCI(7|0)/SCSI(2,0)
```

Table 16-1. SCSI Device Path Examples

Byte Offset	Byte Length	Data	Description
0x00	0x01	0x02	**Generic Device Path Header** – Type ACPI Device Path
0x01	0x01	0x01	Sub type – ACPI Device Path
0x02	0x02	0x0C	Length – 0x0C bytes
0x04	0x04	0x41D0, 0x0A03	_HID PNP0A03 – 0x41D0 represents a compressed string ‘PNP’ and is in the low-order bytes.
0x08	0x04	0x0000	_UID
0x0C	0x01	0x01	**Generic Device Path Header** – Type Hardware Device Path
0x0D	0x01	0x01	Sub type – PCI
0x0E	0x02	0x06	Length – 0x06 bytes
0x10	0x01	0x07	PCI Function
0x11	0x01	0x00	PCI Device
0x12	0x01	0x03	**Generic Device Path Header** – Type Message Device Path
0x13	0x01	0x02	Sub type – SCSI
0x14	0x02	0x08	Length – 0x08 bytes
0x16	0x02	0x0002	Target ID on the SCSI bus, PUN
0x18	0x02	0x0000	Logical Unit Number, LUN
0x1A	0x01	0xFF	**Generic Device Path Header** – Type End of Hardware Device Path
0x1B	0x01	0xFF	Sub type – End of Entire Device Path
0x1C	0x02	0x04	Length – 0x04 bytes

16.5.2 SCSI Device on a Multichannel PCI Controller Example

Table 16-2 shows an example device path for a SCSI device on a multichannel PCI controller. In this example, physical SCSI device 0 is attached on physical channel 1 of a multifunction PCI SCSI

controller. Physical channel 0 is accessed through PCI function #0, and physical channel 1 is accessed through PCI function #1. The following are the device paths for this SCSI device:

ACPI(PNP0A03,1)/PCI(7|1)/SCSI(0,0)

ACPI(PNP0A03,1)/PCI(7|0) Access to controller on physical channel 0

ACPI(PNP0A03,1)/PCI(7|0)/SCSI(0,0)
Access to drive 0 on bus 0 on physical channel 0

ACPI(PNP0A03,1)/PCI(7|1) Access to controller to physical channel 1

ACPI(PNP0A03,1)/PCI(7|1)/SCSI(0,2)
Access to drive 2 on bus 0 on physical channel 1

Table 16-2. SCSI Device on a Multichannel PCI Controller Examples

Byte Offset	Byte Length	Data	Description
0x00	0x01	0x02	**Generic Device Path Header** – Type ACPI Device Path
0x01	0x01	0x01	Sub type – ACPI Device Path
0x02	0x02	0x0C	Length – 0x0C bytes
0x04	0x04	0x41D0, 0x0A03	_HID PNP0A03 – 0x41D0 represents a compressed string 'PNP' and is in the low-order bytes.
0x08	0x04	0x0001	_UID
0x0C	0x01	0x01	**Generic Device Path Header** – Type Hardware Device Path
0x0D	0x01	0x01	Sub type – PCI
0x0E	0x02	0x06	Length – 0x06 bytes
0x10	0x01	0x07	PCI Function
0x11	0x01	0x01	PCI Device
0x12	0x01	0x03	**Generic Device Path Header** – Type Message Device Path
0x13	0x01	0x02	Sub type – SCSI
0x14	0x02	0x08	Length – 0x08 bytes
0x15	0x00	0x0000	Target ID on the SCSI bus, PUN
0x18	0x00	0x0000	Logical Unit Number, LUN
0x1A	0x01	0xFF	**Generic Device Path Header** – Type End of Hardware Device Path
0x1B	0x01	0xFF	Sub type – End of Entire Device Path
0x1C	0x02	0x04	Length – 0x04 bytes
0x00	0x01	0x02	**Generic Device Path Header** – Type ACPI Device Path

16.5.3 ATAPI Device Path Example

Table 16-3 shows an example device path for an ATAPI device on a desktop platform. This ATAPI device is connected to the IDE bus on the primary channel and is configured as the master device on the channel. The IDE bus is generated by the IDE controller, which is a PCI device. It is located at PCI device number 0x1F and PCI function 0x01 and is directly attached to a PCI root bridge.

This sample device path consists of an ACPI device path node, a PCI device path node, an ATAPI node, and a device path end structure. The _HID and _UID must match the ACPI table description of the PCI root bridge. The following is the shorthand notation for this device path:

```
ACPI(PNP0A03,0)/PCI(7|0)/ATAPI(Primary,Master)
```

Table 16-3. ATAPI Device Path Examples

Byte Offset	Byte Length	Data	Description
0x00	0x01	0x02	**Generic Device Path Header** – Type ACPI Device Path
0x01	0x01	0x01	Sub type – ACPI Device Path
0x02	0x02	0x0C	Length – 0x0C bytes
0x04	0x04	0x41D0, 0x0A03	_HID PNP0A03 – 0x41D0 represents a compressed string 'PNP' and is in the low order bytes.
0x08	0x04	0x0000	_UID
0x0C	0x01	0x01	**Generic Device Path Header** – Type Hardware Device Path
0x0D	0x01	0x01	Sub type – PCI
0x0E	0x02	0x06	Length – 0x06 bytes
0x10	0x01	0x07	PCI Function
0x11	0x01	0x00	PCI Device
0x12	0x01	0x03	**Generic Device Path Header** – Type Message Device Path
0x13	0x01	0x01	Sub type – ATAPI
0x14	0x02	0x08	Length – 0x08 bytes
0x16	0x01	0x00	PrimarySecondary – Set to zero for primary or one for secondary.
0x17	0x01	0x00	SlaveMaster – Set to zero for master or one for slave.
0x18	0x02	0x0000	Logical Unit Number, LUN.
0x1A	0x01	0xFF	**Generic Device Path Header** – Type End of Hardware Device Path
0x1B	0x01	0xFF	Sub type – End of Entire Device Path
0x1C	0x02	0x04	Length – 0x04 bytes

16.5.4 Fibre Channel Device Path Example

Table 16-4 shows an example device path for a SCSI device that is connected to a Fibre Channel port on a desktop platform. The Fibre Channel port is a PCI device that is located at PCI device number 0x08 and PCI function 0x00 and is directly attached to a PCI root bridge. The Fibre Channel port is addressed by the World Wide Number and is assigned as X (X is a 64-bit value); the SCSI device's LUN is 0.

This sample device path consists of an ACPI device path node, a PCI device path node, a Fibre Channel device path node, and a device path end structure. The _HID and _UID must match the ACPI table description of the PCI root bridge. The following is the shorthand notation for this device path:

```
ACPI(PNP0A03,0)/PCI(8|0)/Fibre(X,0)
```

Table 16-4. Fibre Channel Device Path Examples

Byte Offset	Byte Length	Data	Description
0x00	0x01	0x02	**Generic Device Path Header** – Type ACPI Device Path
0x01	0x01	0x01	Sub type – ACPI Device Path
0x02	0x02	0x0C	Length – 0x0C bytes
0x04	0x04	0x41D0, 0x0A03	_HID PNP0A03 – 0x41D0 represents a compressed string 'PNP' and is in the low order bytes.
0x08	0x04	0x0000	_UID
0x0C	0x01	0x01	**Generic Device Path Header** – Type Hardware Device Path
0x0D	0x01	0x01	Sub type – PCI
0x0E	0x02	0x06	Length – 0x06 bytes
0x10	0x01	0x08	PCI Function
0x11	0x01	0x00	PCI Device
0x12	0x01	0x03	**Generic Device Path Header** – Type Message Device Path
0x13	0x01	0x02	Sub type – Fibre Channel
0x14	0x02	0x24	Length – 0x24 bytes
0x16	0x04	0x00	Reserved
0x1A	0x08	X	Fibre Channel World Wide Number
0x22	0x08	0x00	Fibre Channel Logical Unit Number.
0x2A	0x01	0xFF	**Generic Device Path Header** – Type End of Hardware Device Path
0x2B	0x01	0xFF	Sub type – End of Entire Device Path
0x2C	0x02	0x04	Length – 0x04 bytes

16.5.5 SCSI Device on a RAID Multichannel Adapter Example

Table 16-5 shows an example device path for a SCSI device on a RAID SCSI host adapter. The PCI SCSI host adapter generates two physical SCSI channels. This SCSI device L is attached on the logical SCSI channel that is generated by a RAID configuration from those two physical channels. The SCSI device L is assigned SCSI ID 0, and its LUN is 0. Figure 16-4 shows the configuration of this adapter.

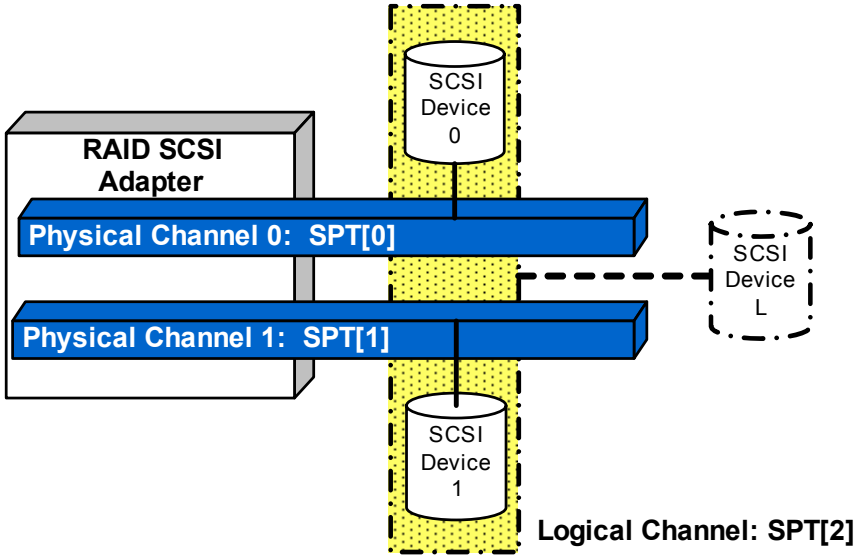

Figure 16-4. Sample RAID SCSI Adapter Configuration

In this example, RAID configures physical SCSI device 0 (which is attached on physical channel 0) and SCSI device 1 (attached on physical channel 1) to SCSI device L, which is attached on the logical SCSI channel. As described in previous sections, only the logical SCSI channel is enumerated at this configuration and only SCSI device L has a valid device path.

Thus, the following is the device path for this SCSI device L:

```
ACPI(PNP0A03,1)/PCI(7|0)/CONTROLLER(2)/SCSI(0,0)
```

Table 16-5. SCSI Device Path Example on RAID Adapter

Byte Offset	Byte Length	Data	Description
0x00	0x01	0x02	**Generic Device Path Header** – Type ACPI Device Path
0x01	0x01	0x01	Sub type – ACPI Device Path
0x02	0x02	0x0C	Length – 0x0C bytes
0x04	0x04	0x41D0, 0x0A03	_HID PNP0A03 – 0x41D0 represents a compressed string ‘PNP’ and is in the low order bytes.
0x08	0x04	0x0001	_UID
0x0C	0x01	0x01	**Generic Device Path Header** – Type Hardware Device Path
0x0D	0x01	0x01	Sub type – PCI
0x0E	0x02	0x06	Length – 0x06 bytes
0x10	0x01	0x07	PCI Function
0x11	0x01	0x00	PCI Device

continued

Table 16-5. SCSI Device Path Example on RAID Adapter (continued)

Byte Offset	Byte Length	Data	Description
0x12	0x01	0x01	**Generic Device Path Header** – Type Hardware Device Path
0x13	0x01	0x05	Sub type – Controller
0x14	0x02	0x08	Length – 0x08 bytes
0x16	0x04	0x02	Controller number
0x1A	0x01	0x03	**Generic Device Path Header** – Type Message Device Path
0x1B	0x01	0x02	Sub type – SCSI
0x1C	0x02	0x08	Length – 0x08 bytes
0x1E	0x00	0x0000	Target ID on the SCSI bus, PUN
0x20	0x00	0x0000	Logical Unit Number, LUN
0x22	0x01	0Xff	**Generic Device Path Header** – Type End of Hardware Device Path
0x23	0x01	0xFF	Sub type – End of Entire Device Path
0x24	0x02	0x04	Length – 0x04 bytes

16.6 Using the SCSI Pass Thru Protocol

If a SCSI driver supports both blocking and nonblocking I/O modes, any client of the SCSI driver can use them to perform SCSI I/O. Example 16-10 demonstrates how to use the SCSI Pass Thru Protocol to perform blocking and nonblocking I/O.

```
EFI_STATUS
ScsiPassThruTests(
  EFI_SCSI_PASS_THRU _PROTOCOL  *EfiSptProtocol
  )
{
  EFI_STATUS                              Status;
  UINT16                                  Target;
  UINT64                                  Lun;
  EFI_SCSI_PASS_THRU_SCSI_REQUEST_PACKET Packet;
  EFI_EVENT                               Event;

  . . .

  //
  // Blocking I/O
  //
  Status = EfiSptProtocol->PassThru (
                                  EfiSptProtocol,
                                  Target,
                                  Lun,
                                  &Packet,
                                  NULL
                                  );
  . . .

  //
  // Non Blocking I/O
  //
```

```
  Status = EfiSptProtocol->PassThru (
                                EfiSptProtocol,
                                Target,
                                Lun,
                                &Packet,
                                &Event
                                );
  . . .

  return Status;
}
```

Example 16-10. Blocking and Nonblocking Modes

17
Driver Optimization Techniques

There are several techniques that can be used to optimize an EFI driver. These techniques can be broken down into the following two categories:

- Techniques to reduce the size of EFI drivers
- Techniques to improve the performance of EFI drivers

Sometimes these techniques complement each other, and sometimes they are at odds with each other. For example, an EFI driver may grow in size to meet a specific performance goal. The driver writer will have to make the appropriate compromises in the selection of these driver optimization techniques.

17.1 Space Optimizations

Table 17-1 lists the techniques that can be used to reduce the size of EFI drivers. By using combinations of all of these techniques, significant size reductions can be realized. The compiler and linker switches that are referenced below are specific to the Microsoft Visual Studio tool chain. Different compilers and linkers may use different switches for equivalent operations.

Table 17-1. Space Optimizations

Technique	Description
EFI Driver Library	The EFI Driver Library should be used to reduce the size of EFI drivers. It is a lightweight library that contains only the functions that most EFI drivers require. The EFI Driver Library consists of three different groups of functions: • General-purpose library functions • String library functions • Print library functions Most EFI drivers use the general-purpose library. EFI drivers that need to manipulate Unicode and ASCII strings use the string library, and EFI drivers that require the ability to generate formatted strings use the print library. Any EFI driver that uses the `DEBUG()` macro will also require the print library.
/Os or /O1 Compiler Switches	Both the /Os and /O1 compiler switches will optimize a C compiler for size. This is an easy way to significantly reduce the size of an EFI driver. Care must be taken when turning on compiler optimizations because C source that works fine with optimizations disabled may stop working with optimizations enabled. They usually stop working because of missing `volatile` declarations on variables and data structures that are shared between normal contexts and raised TPL contexts. Also, because the EFI driver is small, it may execute faster. If there are any speed paths in an EFI driver that will cause problems if the EFI driver executes faster, then these switches may expose those speed paths. These same speed paths will also show up as faster processors are used, so it is good to find these speed paths early.

continued

Table 17-1. Space Optimizations (continued)

Technique	Description
/OPT:REF Linker Switch	This linker switch removes unused functions and variables from the executable image, including functions and variables in the EFI driver and the libraries against which the EFI driver is being linked. The combination of using the EFI Driver Library with this linker switch can significantly reduce the size of an EFI driver executable. Also, `DEBUG()` macros are removed when a production build is performed, so using this linker switch will remove the print library from the executable image.
EFI Compression	If an EFI driver is going to be stored in a PCI option ROM, then the EFI compression algorithm can be used to further reduce the size of an EFI driver. The build utility `EfiRom` has built-in support for compressing EFI images, and the PCI bus driver has built-in support for decompressing EFI drivers stored in PCI option ROMs. The average compression ratio on IA-32 is 2.3, and the average compression ratio on the Itanium processor is 2.8.
EFI Byte Code Images	If an EFI driver is going to be stored in a PCI option ROM and the PCI option ROM must support both IA-32 and Itanium-based platforms or just Itanium-based platforms, then EFI Byte Code (EBC) executables should be considered. EBC executables are portable between IA-32 and Itanium processors. This portability means that only a single EFI driver image is required to support both IA-32 and Itanium-based platforms. Also, the EBC executables are significantly smaller than images for the Itanium processor, so there are advantages to using this format for EFI drivers that are targeted only at Itanium-based platforms. In addition, using *EFI Compression* (see above) can reduce the EBC executables even further.

17.2 Speed Optimizations

Table 17-2 lists the techniques that can be used to improve the performance of EFI drivers. By using combinations of all of these techniques, significant performance enhancements can be realized.

Table 17-2. Speed Optimizations

Technique	Description
/Ot or /O1 Compiler Switches	The /Ot switch optimizes for code speed, and the /O1 compiler switch optimizes a C compiler for size. This technique is an easy way to reduce the execution time and significantly reduce the size of an EFI driver. Care must be taken when turning on compiler optimizations because C source that works fine with optimizations disabled may stop working with optimizations enabled. They usually stop working because of missing `volatile` declarations on variables and data structures that are shared between normal contexts and raised TPL contexts. Also, because the EFI driver is small, it may execute faster. If there are any speed paths in an EFI driver that will cause problems if the EFI driver executes faster, then these switches may expose those speed paths. These same speed paths will also show up as faster processors are used, so it is good to find these speed paths early.

continued

Table 17-2. Speed Optimizations (continued)

Technique	Description
/O2 Compiler Switch	The /O2 compiler switch will optimize a C compiler for speed. This technique is an easy way to improve the performance of an EFI driver. Care must be taken when turning on compiler optimizations because C source that works fine with optimizations disabled may stop working with optimizations enabled. They usually stop working because of missing `volatile` declarations on variables and data structures that are shared between normal contexts and raised TPL contexts. If there are any speed paths in an EFI driver that will cause problems if the EFI driver executes faster, then this switch may expose those speed paths. These same speed paths will also show up as faster processors are used, so it is good to find these speed paths early. The use of the /O2 compiler switch may actually increase the size of an EFI driver.
EFI Services	Whenever possible, use EFI Boot Services, EFI Runtime Services, and the protocol services provided by other EFI drivers. The EFI Boot Services and EFI Runtime Services will likely be native calls that have been optimized for the platform, so there will always be a performance advantage for using these services. Some protocol services might be native, and other protocol services might be EBC images. Either way, if all EFI drivers assume that external protocol services are native, then the combination of EFI drivers and EFI services will result in more efficient execution.
PCI I/O Protocol	If an EFI driver is a PCI driver, then it should take advantage of all the PCI I/O Protocol services to improve the EFI driver's performance. This approach means that all register accesses should be performed at the largest possible size. For example, perform a single 32-bit read instead of multiple 8-bit reads. Also, take advantage of the read/write multiple, FIFO, and fill modes of the `Io()`, `Mem()`, and `Pci()` services.

17.2.1 CopyMem() and SetMem() Operations

Example 17-1 shows examples of how **gBS->CopyMem()** and **gBS->SetMem()** should be used to improve the performance of an EFI driver. These techniques apply to arrays, structures, or allocated buffers.

```
typedef struct {
  UINT8   First;
  UINT32  Second;
} MY_STRUCTURE;

UINTN         Index;
UINT8         A[100];
UINT8         B[100];
MY_STRUCTURE  MyStructureA;
MY_STRUCTURE  MyStructureB;

//
// Using a loop is slow or structure assignments is slow
//
for (Index = 0; Index < 100; Index++) {
  A[Index] = B[Index];
}
MyStructueA = MyStructureB;

//
// Using the optimized CopyMem() Boot Services is fast
//
gBS->CopyMem((VOID *)A, (VOID *)B, 100);
gBS->CopyMem((VOID *)MyStructureA, (VOID *)MyStructureB, sizeof (MY_STRUCTURE));

//
// Using a loop or individual assignment statements is slow
//
for (Index = 0; Index < 100; Index++) {
  A[Index] = 0;
}
MyStructureA.First = 0;
MyStructureA.Second = 0;

//
// Using the optimized SetMem() Boot Service is fast.
//
gBS->SetMem((VOID *)A, 100, 0);
gBS->SetMem((VOID *)&MyStructureA, sizeof (MY_STRUCTURE), 0);
```

Example 17-1. CopyMem() and SetMem() Speed Optimizations

17.2.2 PCI I/O Fill Operations

Example 17-2 shows an example of filling a video frame buffer with zeros on a PCI video controller. The frame buffer is 1 MB of memory-mapped I/O that is accessed through BAR #0 of the PCI video controller. The following four methods of performing this operation are shown, from slowest to fastest:

- **Example 1:** Uses a loop to write to the frame buffer 8 bits at a time.
- **Example 2:** Writes to the frame buffer 32 bits at a time. This second example is better and should provide a 4X increase in performance.
- **Example 3:** Performs a single call to the PCI I/O Protocol to fill the entire frame buffer. This example is even better yet. If the code in Example 17-2 is compiled with an EBC compiler, then the third example will have very significant performance increase over the first two examples.
- **Example 4:** Same as the third example, except the frame buffer is filled 32 bits at a time. This method will have about a 4X increase in performance from the third example.

Two methods that can significantly increase the performance of an EFI driver are taking advantage of the fill operations to eliminate loops and writing to a PCI controller at the largest possible size.

```
EFI_PCI_IO_PROTOCOL  *PciIo;
UINT32               Color32;
UINTN                Count;
UINTN                Index;

//
// This is the slowest method. It performs 0x100000 calls through PCI I/O and
// writes to the frame buffer 8 bits at a time.
//
Color8 = 0;
for (Index = 0; Index < 0x100000; Index++) {
  Status = PciIo->Mem.Write (
                        PciIo,
                        EfiPciIoWidthUint8,        // Width
                        0,                         // Bar Index
                        Index,                     // Offset
                        1,                         // Count
                        &Color8                    // Value
                        );
}

//
// This is a little better. It performs 0x100000/4 calls through PCI I/O and
// writes to the frame buffer 32 bits at a time.
//
Color32 = 0;
for (Index = 0; Index < 0x100000; Index += 4) {
  Status = PciIo->Mem.Write (
                        PciIo,
                        EfiPciIoWidthUint32,       // Width
                        0,                         // Bar Index
                        Index,                     // Offset
                        1,                         // Count
                        &Color32                   // Value
                        );
}
```

```
//
// This is much better. It performs 1 call to PCI I/O, but it is writing the
// frame buffer 8 bits at a time.
//
Color8 = 0;
Count = 0x100000;
Status = PciIo->Mem.Write (
                      PciIo,
                      EfiPciIoWidthFillUint8,    // Width
                      0,                         // Bar Index
                      0,                         // Offset
                      Count,                     // Count
                      &Color8                    // Value
                      );

//
// This is the best method. It performs 1 call to PCI I/O, and it is writing
// the frame buffer 32 bits at a time.
//
Color32 = 0;
Count = 0x100000 / sizeof (UINT32);
Status = PciIo->Mem.Write (
                      PciIo,
                      EfiPciIoWidthFillUint32,   // Width
                      0,                         // Bar Index
                      0,                         // Offset
                      Count,                     // Count
                      &Color32                   // Value
                      );
```

Example 17-2. Speed Optimizations Using PCI I/O Fill Operations

17.2.3 PCI I/O FIFO Operations

Example 17-3 shows an example of writing a sector to an IDE controller. The IDE controller uses a single 16-bit I/O port as a FIFO for reading and writing sector data. The first example calls the PCI I/O Protocol 256 times to write the sector. The second example is much better, because it calls the PCI I/O Protocol only once, which will provide a significant performance increase if this example is compiled with an EBC compiler. This example would apply equally to FIFO read operations.

```
EFI_PCI_IO_PROTOCOL  *PciIo;
UINT16               Buffer[256];
UINTN                Index;

//
// This is the slowest method. It performs 256 PCI I/O calls to write 256
// 16-bit values to the IDE controller.
//
for (Index = 0; Index < 256; Index++) {
  Status = PciIo->Io.Write (
                       PciIo,
                       EfiPciIoWidthUint16,
                       EFI_PCI_IO_PASS_THROUGH_BAR,
                       0x1F0,
                       1,
                       Buffer[Index]
                       );
}
```

```
//
// This is the fastest method. It uses a loop to write 256 16-bit values to
// the IDE controller.
//
Status = PciIo->Io.Write (
                     PciIo,
                     EfiPciIoWidthFifoUint16,
                     EFI_PCI_IO_PASS_THROUGH_BAR,
                     0x1F0,
                     256,
                     Buffer
                     );
```

Example 17-3. Speed Optimizations Using PCI I/O FIFO Operations

17.2.4 PCI I/O CopyMem() Operations

Example 17-4 shows an example of scrolling a frame buffer by one scan line on a PCI video controller. Just like the **`gBS->CopyMem()`** Boot Service, the **`CopyMem()`** service in the PCI I/O Protocol should be used whenever it can eliminate loops. The frame buffer is 1 MB of memory-mapped I/O that is accessed through BAR #0 of the PCI video controller, and the screen is 800 pixels wide with 32 bits per pixel. The code below shows the following two methods, from slowest to fastest:

- **Example 1:** Uses a loop to read and write every 32-bit pixel into the variable *`Value`*.
- **Example 2:** Uses a single call to **`CopyMem()`** to perform the exact same function. This second example would be significantly faster if it was compiled with an EBC compiler.

```
EFI_PCI_IO_PROTOCOL  *PciIo;
UINTN                Index;
UINT32               Value;
UINTN                ScanLineWidth;

//
// This is the slowest method. It performs almost 0x100000/4 read and write
// accesses through the PCI I/O protocol.
//
ScanLineWidth = 800 * 4;
for (Index = ScanLineWidth; Index < 0x100000; Index += 4) {
  Status = PciIo->Mem.Read (
                        PciIo,
                        EfiPciIoWidthUint32,    // Width
                        0,                      // Bar Index
                        Index,                  // Offset
                        1,                      // Count
                        &Value                  // Value
                        );
  Status = PciIo->Mem.Write (
                        PciIo,
                        EfiPciIoWidthUint32,    // Width
                        0,                      // Bar Index
                        Index - ScanLineWidth,  // Offset
                        1,                      // Count
                        &Value                  // Value
                        );
}
```

```
//
// This is the fastest method. It makes a single call to CopyMem().
//
Status = PciIo->CopyMem (
                  PciIo,
                  EfiPciIoWidthUint32,     // Width
                  0,                       // Destination Bar Index
                  0,                       // Destination Offset
                  0,                       // Source Bar Index
                  ScanLineWidth,           // Source Offset
                  0x100000/4               // Count
                  );
```

Example 17-4. Speed Optimizations Using the PCI I/O CopyMem() Service

17.2.5 PCI Configuration Header Operations

Example 17-5 shows the following three examples for reading the PCI configuration header from a PCI controller, from slowest to fastest:

- **Example 1:** Uses a loop to read the header 8 bits at a time.
- **Example 2:** Uses a single call to read the entire header 8 bits at a time.
- **Example 3:** Makes a single call to read the header 32 bits at a time.

```
EFI_PCI_IO_PROTOCOL  *PciIo;
PCI_TYPE00           Pci;
UINTN                Index;

//
// Loop reading the 64-byte PCI configuration header 8 bits at a time
//
for (Index = 0; Index < sizeof (Pci); Index++) {
  Status = PciIo->Pci.Read (
                        PciIo,
                        EfiPciIoWidthUint8,      // Width
                        Index,                   // Offset
                        1,                       // Count
                        (UINT8 *)(&Pci) + Index  // Value
                        );
}

//
// This is a faster method that removes the loop and reads 8 bits at a time.
//
Status = PciIo->Pci.Read (
                      PciIo,
                      EfiPciIoWidthUint8,      // Width
                      0,                       // Offset
                      sizeof (Pci),            // Count
                      &Pci                     // Value
                      );

//
// This is the fastest method that makes a single call to PCI I/O and reads the
// PCI configuration header 32 bits at a time.
//
Status = PciIo->Pci.Read (
```

```
                        PciIo,
                        EfiPciIoWidthUint32,                // Width
                        0,                                  // Offset
                        sizeof (Pci) / sizeof (UINT32),     // Count
                        &Pci                                // Value
                        );
```

Example 17-5. Speed Optimizations for PCI Configuration Cycles

17.2.6 PCI I/O Read/Write Multiple Operations

Example 17-6 shows an example of writing to a PCI memory-mapped I/O buffer. This example shows a full-screen bitmap being written to a frame buffer on a PCI video controller. The frame buffer is 1 MB of memory-mapped I/O that is accessed through BAR #0 of the PCI video controller. The code below shows the following two methods, from slowest to fastest:

- **Example 1:** Uses a loop to perform 0x100000 8-bit write operations. This method would be very slow with an EFI driver compiled for EBC.
- **Example 2:** Performs the same operation using a single call to the PCI I/O Protocol, and the write operations are performed 32 bits at a time.

The examples show here apply equally well to reading a bitmap from the frame buffer of a PCI video controller using the **`PciIo->Mem.Read()`** function.

```
EFI_PCI_IO_PROTOCOL  *PciIo;
UINTN                Index;
UINT8                Bitmap[0x100000];

//
// Loop writing a 1 MB bitmap to the frame buffer 8 bits at a time.
//
for (Index = 0; Index < sizeof (BitMap); Index++) {
  Status = PciIo->Mem.Write (
                        PciIo,
                        EfiPciIoWidthUint8,       // Width
                        0,                        // BarIndex
                        Index,                    // Offset
                        1,                        // Count
                        &BitMap[Index]            // Value
                        );
}

//
// This is a faster method that removes the loop and writes 32 bits at a time.
//
Status = PciIo->Mem.Write (
                        PciIo,
                        EfiPciIoWidthUint32,                  // Width
                        0,                                    // BarIndex
                        0,                                    // Offset
                        sizeof (BitMap) / sizeof (UINT32),    // Count
                        BitMap                                // Value
                        );
```

Example 17-6. Speed Optimizations for Multiple Read/Write Operations

17.2.7 PCI I/O Polling Operations

Example 17-7 shows a common example of polling an I/O port for a status bit to change. This poll is usually done when an EFI driver is waiting for the hardware to complete an operation, and the completion status is indicated by a bit changing state in an I/O port or a memory-mapped I/O port. The example shown below polls offset 0x20 in BAR #1 for bit 0 to change from a 0 to a 1. The code below shows the following two methods, from slowest to fastest:

- **Example 1:** Uses a loop with 10 µS stalls to wait up to 1 second for the bit to change in value.
- **Example 2:** Same as example 1, but it makes only a single call to the PCI I/O Protocol to perform the same operation. The second example will execute more efficiently for an EFI driver compiled for EBC.

The **PollIo()** and **PollMem()** functions in the PCI I/O Protocol are very flexible, and they can simplify the operation of polling for bits to change state in status registers.

```
EFI_PCI_IO_PROTOCOL  *PciIo;
UINTN                TimeOut;
UINT8                Result8;
UINT64               Result64;

//
// Loop for up to 1 second waiting for Bit #0 in register 0x20 of BAR #1 to
// become a 1.
//
TimeOut = 0;
do {
  Status = PciIo->Mem.Read (
                        PciIo,
                        EfiPciIoWidthUint8,      // Width
                        1,                       // BarIndex
                        0x20,                    // Offset
                        1,                       // Count
                        &Result8                 // Value
                        );
  if ((Result8 & 0x01) == 0x01) {
    return EFI_SUCCESS;
  }
  gBS->Stall (10);
  TimeOut = TimeOut + 10;
} while (TimeOut < 1000000);
return EFI_TIMEOUT;

//
// Call PollIo() to poll for Bit #0 in register 0x20 of Bar #1 to be set to a 1.
//
Status = PciIo->Pci.PollIo (
                      PciIo,
                      EfiPciIoWidthUint8,      // Width
                      1,                       // BarIndex
                      0x20,                    // Offset
                      0x01,                    // Mask
                      0x01,                    // Value
                      10000000,                // Poll for 1 second
                      &Result64                // Result
                      );
```

Example 17-7. Speed Optimizations for Polled I/O Operations

18
Itanium Architecture Porting Considerations

When writing an EFI driver, there are several steps that can be taken to ensure that the driver will function properly on an Itanium-based platform. Typically, EFI drivers are initially developed for an IA-32 platform and then will be ported to an Itanium-based platform. If an EFI driver contains IA-32 assembly language sources, then those sources must be converted to C or assembly language sources for the Itanium processor. In general, it is always better to write EFI drivers in C so the driver will be as portable as possible. The guidelines listed in this chapter will help improve the portability of an EFI driver and it specifically addresses the pitfalls that may be encountered when an EFI driver is ported to an Itanium processor.

18.1 Alignment Faults

The single largest issue with EFI drivers for Itanium-based platforms is alignment. An IA-32 processor will allow any sized transaction on any byte boundary. The Itanium processor allows transactions to be performed only on natural boundaries. This requirement means that a 64-bit read or write transaction must begin on an 8-byte boundary, a 32-bit read or write transaction must begin on a 4-byte boundary, and a 16-bit read or write transaction must begin on a 2-byte boundary. In most cases, the driver writer does not need to worry about this issue because the C compiler will guarantee that accessing global variables, local variables, and fields of data structures will not cause an alignment fault. The only cases in which C code can generate an alignment fault are when a pointer is cast from one type to another or when packed data structures are used. Alignment faults can also be generated from assembly language, but it is the assembly programmer's responsibility to ensure alignment faults are not generated.

Example 18-1 shows an example that will generate an alignment fault on an Itanium processor. The first read access through *SmallValuePointer* is aligned because *LargeValue* is on a 64-bit boundary. However, the second read access though *SmallValuePointer* will generate an alignment fault because *SmallValuePointer* is not on a 32-bit boundary. The problem is that an 8-bit pointer was cast to a 32-bit pointer. Whenever a cast is made from a pointer to a smaller data type to a pointer to a larger data type, there is a chance that the pointer to the larger data type will be unaligned.

```
UINT64  LargeValue;
UINT32  *SmallValuePointer;
UINT32  SmallValue;

SmallValuePointer  = (UINT32 *)&LargeValue;
SmallValue         = *SmallValuePointer;                    // Works
SmallValuePointer  = (UINT32 *)((UINT8 *)&LargeValue + 1);
SmallValue         = *SmallValuePointer;                    // Faults
```

Example 18-1. Pointer-Cast Alignment Fault

Example 18-2 shows the same example as Example 18-1, but it has been modified to prevent the alignment fault. The second read access through *SmallValuePointer* is replaced with a macro that treats the 32-bit value as an array of bytes. The individual bytes are read and combined into a 32-bit value. The generated object code is larger and slower, but it is functional on both IA-32 and Itanium processors.

```
#define UNPACK_UINT32(a)  (UINT32)( (((UINT8 *) a)[0] <<  0) |     \
                                    (((UINT8 *) a)[1] <<  8) |     \
                                    (((UINT8 *) a)[2] << 16) |     \
                                    (((UINT8 *) a)[3] << 24) )

UINT64  LargeValue;
UINT32  *SmallValuePointer;
UINT32  SmallValue;

SmallValuePointer  = (UINT32 *)&LargeValue;
SmallValue         = *SmallValuePointer;                    // Works
SmallValuePointer  = (UINT32 *)((UINT8 *)&LargeValue + 1);
SmallValue         = UNPACK_UINT32 (SmallValuePointer);     // Works
```

Example 18-2. Corrected Pointer-Cast Alignment Fault

Example 18-3 shows another example that will generate an alignment fault on an Itanium processor. The first read access from *MyStructure.First* will always work because the 8-bit value is always aligned. However, the second read access from *MyStructure.Second* will always fail because the 32-bit value will never be aligned on a 4-byte boundary.

```
#pragma pack(1)
typedef struct {
  UINT8   First;
  UINT32  Second;
} MY_STRUCTURE;
#pragma pack()

MY_STRUCTURE  MyStructure;
UINT8         FirstValue;
UINT32        SecondValue;

FirstValue  = MyStructure.First;   // Works
SecondValue = MyStructure.Second;  // Faults
```

Example 18-3. Packed Structure Alignment Fault

Example 18-4 shows the same example as Example 18-3, but it has been modified to prevent the alignment fault. The second read access from *MyStructure.Second* is replaced with a macro that treats the 32-bit value as an array of bytes. The individual bytes are read and combined into a 32-bit value. The generated object code is larger and slower, but it is functional on both IA-32 and Itanium processors.

```
#define UNPACK_UINT32(a)  (UINT32)( (((UINT8 *) a)[0] <<  0) |     \
                                    (((UINT8 *) a)[1] <<  8) |     \
                                    (((UINT8 *) a)[2] << 16) |     \
                                    (((UINT8 *) a)[3] << 24) )

#pragma pack(1)
typedef struct {
  UINT8   First;
  UINT32  Second;
```

```
} MY_STRUCTURE;
#pragma pack()

MY_STRUCTURE  MyStructure;
UINT8         FirstValue;
UINT32        SecondValue;

FirstValue  = MyStructure.First;                  // Works
SecondValue = UNPACK_UINT32(&MyStructure.Second);  // Works
```

Example 18-4. Corrected Packed Structure Alignment Fault

If a data structure is copied from one location to another, then both the source and the destination pointers for the copy operation should be aligned on a 64-bit boundary for Itanium platforms. The **`gBS->CopyMem()`** service can handle unaligned copy operations, so an alignment fault will not be generated by the copy operation itself. However, if the fields of the data structure at the destination location are accessed, they may generate alignment faults if the destination address is not aligned on a 64-bit boundary. There are cases where an aligned structure may be copied to an unaligned destination, but the fields of the destination buffer must not be accessed after the copy operation is completed. An example of this case is when a packed structure is being built that will be stored on disk or transmitted on a network.

In some cases, it may be necessary to copy a data structure from an unaligned source location to an aligned destination location so that the fields of the data structure can be accessed without generating an alignment fault. The following are two examples of this scenario:

- Parsing EFI device path nodes
- Parsing network packets

The device path nodes in an EFI device path are packed together so they will take up as little space as possible when they are stored in environment variables such as *`ConIn`*, *`ConOut`*, *`StdErr`*, *`Boot####`*, and *`Driver####`*. As a result, individual device path nodes may not be aligned on a 64-bit boundary. EFI device paths or device paths nodes can be passed around as opaque data structures, but whenever the fields of a device path node need to be accessed, the device path node must be copied to a location that is guaranteed to be on a 64-bit boundary. Likewise, network packets are packed so they take up as little space as possible, so as each layer of a network packet is examined, it may need to be copied to a 64-bit aligned location before the individual fields of the packet are examined.

Example 18-5 shows an example of a function that parses an EFI device path and extracts the 32-bit HID and UID from an ACPI device path node. This example will generate an alignment fault if *`DevicePath`* is not aligned on a 32-bit boundary.

```
VOID
GetAcpiHidUid (
  EFI_DEVICE_PATH_PROTOCOL  *DevicePath,
  UINT32                    *Hid,
  UINT32                    *Uid
  )

{
  ACPI_HID_DEVICE_PATH  *AcpiDevicePath;

  AcpiDevicePath = (ACPI_HID_DEVICE_PATH *)DevicePath;
```

```
  *Hid = AcpiDevicePath.HID;            // May fault
  *Uid = AcpiDevicePath.UID;            // May fault
}
```

Example 18-5. EFI Device Path Node Alignment Fault

Example 18-6 shows the corrected version of Example 18-5. Because the alignment of *DevicePath* cannot be guaranteed, the solution is to copy the ACPI device path node from *DevicePath* into an ACPI device path node structure that is declared as the local variable *AcpiDevicePath*. A structure declared as a local variable is guaranteed to be on a 64-bit boundary on Itanium platforms. The fields of the ACPI device path node can then be safely accessed without generating an alignment fault.

```
VOID
GetAcpiHidUid (
  EFI_DEVICE_PATH_PROTOCOL  *DevicePath,
  UINT32                    *Hid,
  UINT32                    *Uid
  )

{
  ACPI_HID_DEVICE_PATH  AcpiDevicePath;

  gBS->CopyMem(&AcpiDevicePath, DevicePath, sizeof (ACPI_HID_DEVICE_PATH));
  *Hid = AcpiDevicePath.HID;            // Guaranteed to work
  *Uid = AcpiDevicePath.UID;            // Guaranteed to work
}
```

Example 18-6. Corrected EFI Device Path Node Alignment Fault

18.2 Accessing a 64-Bit BAR in a PCI Configuration Header

Another source of alignment faults is when 64-bit BAR values are accessed in a PCI configuration header. A PCI configuration header has room for up to six 32-bit BAR values or three 64-bit BAR values. A PCI configuration header may also contain a mix of both 32-bit BAR values and 64-bit BAR values. All 32-bit BAR values are guaranteed to be on a 32-bit boundary. However, 64-bit BAR values may be on a 32-bit boundary or a 64-bit boundary. As a result, every time a 64-bit BAR value is accessed, it must be assumed to be on a 32-bit boundary to guarantee that an alignment fault will not be generated. The following are a couple of methods that may be used to prevent an alignment fault when a 64-bit BAR value is extracted from a PCI configuration header:

- **Method 1:** Use **`gBS->CopyMem()`** to transfer the BAR contents into a 64-bit aligned location.
- **Method 2:** Collect the two 32-bit values that compose the 64-bit BAR and combine them into a 64-bit value.

Example 18-7 below shows the incorrect method of extracting a 64-bit BAR from a PCI configuration header and two correct methods.

```
UINT64
Get64BitBarValue (
  PCI_TYPE00  *PciConfigurationHeader,
  UINTN       BarOffset
  )
```

```
{
  UINT64  *BarPointer64;
  UINT32  *BarPointer32;
  UINT64  BarValue;

  BarPointer64 = (UINT64 *)((UINT8 *)PciConfigurationHeader + BarOffset);
  BarPointer32 = (UINT32 *)((UINT8 *)PciConfigurationHeader + BarOffset);

  //
  // Wrong. May cause an alignment fault.
  //
  BarValue = *BarPointer64;

  //
  // Correct. Guaranteed not to generate an alignment fault.
  //
  gBS->CopyMem (&BarValue, BarPointer64, sizeof (UINT64));

  //
  // Correct. Guaranteed not to generate an alignment fault.
  //
  BarValue = (UINT64)(*BarPointer32 | LShiftU64 (*(BarPointer32 + 1), 32));

  return BarValue;
}
```

Example 18-7. Accessing a 64-Bit BAR in a PCI Configuration Header

18.3 Assignment and Comparison Operators

There are issues if a data value is cast from a larger size to a smaller size. In these cases, the upper bits of the larger values are stripped. In general, this stripping will cause a compiler warning, so these are easy issues to catch. However, there are a few cases where everything compiles free of errors and warnings on IA-32 and generates errors or warnings on the Itanium processor. The only way to guarantee that these errors are caught early is to compile for both IA-32 and Itanium processors during the entire development process. When one of these warnings is generated by an Intel® Itanium® compiler, the warning can be eliminated by explicitly casting the larger data type to the smaller data type. However, the developer needs to make sure that this casting is the right solution, because the upper bits of the larger data value will be stripped.

Example 18-8 shows several examples that will generate a warning and how to eliminate the warning with an explicit cast. The last example is the most interesting one because it does not generate any warnings on IA-32, but it will on Itanium architecture. This difference is because a **UINTN** on IA-32 is identical to **UINT32**, but **UINTN** on Itanium architecture is identical to a **UINT64**.

```
UINT8    Value8;
UINT16   Value16;
UINT32   Value32;
UINT64   Value64;
UINTN    ValueN;

Value8  = Value16;           // Warning generated on IA-32 and Itanium
Value8  = (UINT8)Value16;    // Works, upper 8 bits stripped
```

```
Value16 = Value8;            // Works

Value8  = Value32;           // Warning generated on IA-32 and Itanium
Value8  = (UINT8)Value32;    // Works, upper 24 bits stripped
Value32 = Value8;            // Works

Value8  = Value64;           // Warning generated on IA-32 and Itanium
Value8  = (UINT8)Value64;    // Works, upper 56 bits stripped
Value64 = Value8;            // Works

Value8  = ValueN;            // Warning generated on IA-32 and Itanium
Value8  = (UINT8)ValueN;     // Works, upper 24 bits stripped on IA-32.
                             // Upper 56 bits stripped on Itanium
ValueN  = Value8;            // Works

Value32 = ValueN;            // Works on IA-32, warning generated on Itanium
Value32 = (UINT32)ValueN;    // Works on IA-32, upper 32 bits stripped on
                             // Itanium
```

Example 18-8. Assignment Operation Warnings

Example 18-9 is very similar to Example 18-8, except the assignments have been replaced with comparison operations. The same issues shown will be generated by all the comparison operators, including **>**, **<**, **>=**, **<=**, **!=**, and **==**. The solution is to cast one of the two operands to be the same as the other operand. The first four cases are the ones that work on IA-32 with no errors or warnings but generate warnings on Itanium architecture. The next four cases resolve the issue by casting the first operand, and the last four cases resolve the issue by casting the second operand. Care must be taken when casting the correct operand, because a cast from a larger data type to a smaller data type will strip the upper bits of the operand. When a cast is performed to **INTN** or **UINTN**, a different number of bits will be stripped for IA-32 and Itanium architecture.

```
UINT64  ValueU64;
UINTN   ValueUN;
INT64   Value64;
INTN    ValueN;

if (ValueU64 == ValueN)   {}  // Works on IA-32, warning generated on Itanium
if (ValueUN  == Value64)  {}  // Works on IA-32, warning generated on Itanium
if (Value64  == ValueUN)  {}  // Works on IA-32, warning generated on Itanium
if (ValueN   == ValueU64) {}  // Works on IA-32, warning generated on Itanium

if ((INTN)ValueU64 == ValueN)   {}  // Works on both IA-32 and Itanium
if ((INT64)ValueUN == Value64)  {}  // Works on both IA-32 and Itanium
if ((UINTN)Value64 == ValueUN)  {}  // Works on both IA-32 and Itanium
if ((UINT64)ValueN == ValueU64) {}  // Works on both IA-32 and Itanium

if (ValueU64 == (UINT64)ValueN) {}  // Works on both IA-32 and Itanium
if (ValueUN  == (UINTN)Value64) {}  // Works on both IA-32 and Itanium
if (Value64  == (INT64)ValueUN) {}  // Works on both IA-32 and Itanium
if (ValueN   == (INTN)ValueU64) {}  // Works on both IA-32 and Itanium
```

Example 18-9. Comparison Operation Warnings

18.4 Casting Pointers

Pointers can be cast from one pointer type to another pointer type. However, pointers should never be cast to a fixed-size data type, and fixed-size data types should never be cast to pointers. The size of a pointer varies between IA-32 and Itanium processors. If any assumptions are made that a pointer to a function or a pointer to a data structure is a 32-bit value, then that code will not run on Itanium-based platforms with physical memory above 4 GB. These issues difficult are to catch, because explicit casts are required to cast a fixed-width type to a pointer or vice versa. Once these explicit type casts are introduced, no compiler warnings or errors will be generated. In fact, the code may execute just fine on IA-32 platforms and Itanium-based platforms with physical memory below 4 GB. The only failing case will be when the code is tested on an Itanium-based system with physical memory above 4 GB. The symptom is typically a processor exception that results in a system hang or reset. Example 18-10 below shows some good and bad examples of casting pointers. The first group is casting pointers to pointers. The second group is casting pointers to fixed width types, and the last group is casting fixed width types to pointers. There is one exception to this rule that applies to IA-32 and Itanium processors. The data types **INTN** and **UINTN** are the exact same size of pointers on both IA-32 and Itanium-based platforms, which means that a pointer can be cast to or from **INTN** or **UINTN** without any adverse side effects. However, ANSI C does not require function pointers to be the same size as data pointers, and function pointers and data pointers are not required to be the same size as **INTN** or **UINTN**. As a result, this exception does not apply to all processors.

```
typedef struct {
  UINT8   First;
  UINT32  Second;
} MY_STRUCTURE;

MY_STRUCTURE  *MyStructure;
UINT8         ValueU8;
UINT16        ValueU16;
UINT32        ValueU32;
UINT64        ValueU64;
UINTN         ValueUN;
INT64         Value64;
INTN          ValueN;
VOID          *Pointer;

//
// Casting pointers to pointers
//
Pointer     = (VOID *)MyStructure;      // Good.
MyStructure = (MY_STRUCTURE *)Pointer;  // Good.

//
// Casting pointers to fixed width types
//
ValueU8  = (UINT8)MyStructure;     // Bad. Strips upper 24 bits on IA-32 and
                                   // upper 56 bits on Itanium.
ValueU16 = (UINT16)MyStructure;    // Bad. Strips upper 16 bits on IA-32 and
                                   // upper 48 bits on Itanium.
ValueU32 = (UINT32)MyStructure;    // Bad. Works on IA-32, but strips upper
                                   // 32 bits on Itanium.
ValueU64 = (UINT64)MyStructure;    // OK. Works on IA-32 and Itanium
Value64  = (INT64)MyStructure;     // OK. Works on IA-32 and Itanium
```

```
ValueUN  = (UINTN)MyStructure;     // Good. Works on IA-32 and Itanium
ValueN   = (INTN)MyStructure;      // Good. Works on IA-32 and Itanium

//
// Casting fixed width types to pointers
//
MyStructure = (MY_STRUCTURE *)ValueU8;   // Bad
MyStructure = (MY_STRUCTURE *)ValueU16;  // Bad
MyStructure = (MY_STRUCTURE *)ValueU32;  // Bad. Works on IA-32, only works on
                                         // Itanium-based platforms with < 4 GB
MyStructure = (MY_STRUCTURE *)ValueU64;  // OK. Works on IA-32 and Itanium
MyStructure = (MY_STRUCTURE *)Value64;   // OK. Works on IA-32 and Itanium
MyStructure = (MY_STRUCTURE *)ValueUN;   // Good. Works on IA-32 and Itanium
MyStructure = (MY_STRUCTURE *)ValueN;    // Good. Works on IA-32 and Itanium
```

Example 18-10. Casting Pointer Examples

18.5 EFI Data Type Sizes

There are a few EFI data types that are different sizes on IA-32 and Itanium architecture, as follows:

- Pointers
- Enumerations
- **INTN**
- **UINTN**

These differing types also mean that any complex types, such as unions and data structures, that are composed of these base types will also have different sizes on IA-32 and Itanium architecture. These differences must be understood whenever the **sizeof()** operator is used. If a union or data structure is required that does not change size between IA-32 and Itanium architecture, see Appendix A for a summary of the EFI data types that are available to all EFI applications and EFI drivers.

18.6 Negative Numbers

Negative numbers are not the same on IA-32 and Itanium processors. Negative numbers are type **INTN**, and **INTN** is a 4-byte container on IA-32 and an 8-byte container on the Itanium processor. For example, **-1** on IA-32 is **0xffffffff**, and **-1** on the Itanium processor is **0xffffffffffffffff**. Care must be taken when assigning or comparing negative numbers. Example 18-11 shows an example that compiles without errors or warnings on both IA-32 and Itanium processors but behaves very differently on IA-32 than it does on the Itanium processor.

```
UINT32  ValueU32;

ValueU32 = 0xffffffff;

if ((INTN)ValueU32 == -1) {
  Print(L"Equal\n");       // This message is printed on IA-32 and not Itanium
} else {
  Print(L"Not Equal\n");   // This message is printed on Itanium and not IA-32
}
```

Example 18-11. Negative Number Example

18.7 Returning Pointers in a Function Parameter

Example 18-12 shows a bad example for casting pointers. The function **MyFunction()** simply returns a 64-bit value in an **OUT** parameter that is assigned from a 32-bit input parameter. There is nothing wrong with **MyFunction()**. The problem is when **MyFunction()** is called. Here, the address of **B**, a 32-bit container, is cast to a pointer to a 64-bit container and passed to **MyFunction()**. **MyFunction()** writes to 64 bits starting at **B**. This location happens to overwrite the value of **B** and the value of **A** in the calling function. The first **Print()** correctly shows the values of **A** and **B**. The second **Print()** shows that **B** was given **A**'s original value, but the contents of **A** were destroyed and overwritten with a 0. The cast from **&B** to a **(UINT64 *)** is the problem here. This code compiles without errors or warnings in both IA-32 and Itanium processors. It executes on IA-32 with these unexpected side effects. It might run on Itanium processors, but it depends on if **&B** is 64-bit aligned or not. There is a 50 percent chance that it will generate an alignment fault on an Itanium processor. If it does not generate an alignment fault, then it will get the same unexpected results that occurred on IA-32. This porting issue is not specific to the Itanium processor. However, because 64-bit quantities are more likely to be used in EFI drivers, this issue is an important one to consider when doing EFI development work.

```
EFI_STATUS
MyFunction (
  IN  UINT32  ValueU32,
  OUT UINT64  *ValueU64
  )

{
  *ValueU64 = (UINT64)ValueU32;
  return EFI_SUCCESS;
}


UINT32  A;
UINT32  B;

A = 0x11112222;
B = 0x33334444;
Print(L"A = %08x  B = %08x\n",A,B);  // Prints "A = 11112222  B = 33334444"
MyFunction (A, (UINT64 *)(&B));
Print(L"A = %08x  B = %08x\n",A,B);  // Prints "A = 00000000  B = 11112222"
```

Example 18-12. Casting OUT Function Parameters

18.8 Array Subscripts

In general, array subscripts should be of type **INTN** or **UINTN**. Using these types will avoid problems if an array subscript is decremented below **0**. If a **UINT32** is used as an array subscript and is decremented below 0, it is decremented to **0xffffffff** on IA-32 and **0x00000000ffffffff** on the Itanium processor. These subscript values are very different. On IA-32, this value is the same indexing element as **-1** of the array. However, on the Itanium processor, this value is the same indexing element as **0xffffffff** of the array. If an **INTN** or **UINTN** is used instead of a **UINT32** for the array subscript, then this problem goes away. When a **UINTN** is decremented below **0**, it is decremented to **0xffffffff** on IA-32 and

0xffffffffffffffff on the Itanium processor. These values are both the same indexing element as **-1** of the array. Example 18-13 below show two examples of array subscripts. The first one works on IA-32 but faults on Itanium processors. The second example is rewritten to work properly on both IA-32 and Itanium processors.

```
UINT32  Index;
CHAR8   Array[] = "ABCDEFGHIJKLIMNOPQRSTUVWXYZ";
CHAR8   *MyArray;

MyArray = &(Array[5]);
Index = 0;
Print(L"Character = %c\n",Array[Index-1]);  // Works on IA-32, faults on Itanium


UINTN   Index;
CHAR8   Array[] = "ABCDEFGHIJKLIMNOPQRSTUVWXYZ";
CHAR8   *MyArray;

MyArray = &(Array[5]);
Index = 0;
Print(L"Character = %c\n",Array[Index-1]);  // Works on IA-32 and Itanium
```

Example 18-13. Array Subscripts Example

18.9 Piecemeal Structure Allocations

Structures should always be allocated using the **sizeof()** operator on the name of the structure. The sum of the sizes of the structure's members should never be used because it does not take into account the padding that the compiler introduces to guarantee alignment. Example 18-14 shows two examples for allocating memory for a structure. The first one is incorrect and the second allocation is correct.

```
typedef struct {
  UINT8   Value8;
  UINT64  Value64;
} MY_STRUCTURE;

MY_STRUCTURE  *MyStructure;

//
// Wrong. This will only allocate 9 bytes, but MyStructure is 16 bytes
//
Status = gBS->AllocatePool (
                EfiBootServicesData,
                sizeof (UINT8) + sizeof (UINT64),
                (VOID **)&MyStructure
                );

//
// Correct. This will allocate 16 bytes for MyStructure.
//
Status = gBS->AllocatePool (
                EfiBootServicesData,
                sizeof (MY_STRUCTURE),
                (VOID **)&MyStructure
                );
```

Example 18-14. Piecemeal Structure Allocation

18.10 Speculation and Floating Point Register Usage

Itanium processors support speculative memory accesses and a large number of floating point registers. EFI drivers that are compiled for Itanium processors must follow the calling conventions defined in the *SAL Specification.* This specification allows only the first 32 floating point registers to be used and defines the amount of speculation support that a platform is required to implement for the preboot environment. These requirements mean that the correct compiler and linker switches must be used to guarantee that these calling conventions are followed. The *EFI Sample Implementation* includes the correct compiler and linker settings for several tool chains for the Itanium processor. These settings may have to be adjusted if a newer or different tool chain is used. Table 18-1 shows the compiler flags for a few different compilers.

Table 18-1. Compiler Flags

Compiler	Optimizations	Description
Intel Itanium Compiler 5.01	Off	/X /Zl /Zi /Od /W3 /QIPF_fr32
	On	/X /Zl /O1 /W3 /QIPF_fr32
Intel® C++ Compiler 7.1 for Windows*	Off	/X /Zi /Zl /Od /W3 /WX /QIPF_fr32
	On	/X /Zl /O1 /W3 /WX /QIPF_fr32

18.11 Memory Ordering

The store model of the processor or the store model for a bus master in an I/O subsystem may be weakly ordered. Weak ordering of processor store cycles has been the source of several difficult bugs in a number of drivers. It is easy to imagine that this issue could be fairly widespread in drivers written for Itanium processors at both the OS and EFI level. See *A Formal Specification of Intel Itanium Processor Family Memory Ordering* for a detailed discussion of this topic, which is also discussed in the *Intel® Itanium® Architecture Software Developer Manuals,* volumes 1–4.

The classic case where strong ordering versus weak ordering produces different results is when there is a memory-based FIFO and a shared bus master "doorbell" register that is shared by all additions to the FIFO. In this common implementation, the driver (producer) formats a new request descriptor and, as its last logical operation, writes the value indicating the entry is valid.

This mechanism becomes a problem if a new request is being added to the FIFO while the bus master is checking the next FIFO entry's valid flag. It is possible for the "last write" issued by the processor (that turns on the valid flag) to be posted to memory before the logically earlier writes that finish initializing the FIFO/request descriptor.

The solution in this case is to ensure that all pending memory writes have been completed before the "valid flag" is enabled. There are two techniques to avoid this problem:

- **Technique 1:** Declare the whole structures with the C language "volatile" attribute. The compiler will ensure that strong ordering is used for all operations in this case.
- **Technique 2:** Use the `MEMORY_FENCE()` macro before setting the valid flag. This macro calls the "`__mf();`" intrinsic, which will ensure that all previous stores are posted. The intrinsic call requires that a `#pragma intrinsic (__mfa)` statement be defined. The

EFI 1.10 Sample Implementation sets up this macro in the **`\EFI1.1\inc\ipf\efibind.h`** header file.

The second solution is typically preferred for readability because the intent is clearer. A volatile declaration tends to hide what was needed, because it is not part of the affected code (because it is off in a structure definition). In addition, using the volatile declaration could impact the driver's performance because all memory transactions to the structure would be strongly ordered (ordered memory transactions are slower).

NOTE

If a driver is executing in the EBC environment, the EBC interpreter ensures that all memory transactions are strongly ordered. Technically, EBC drivers do not need to use the **`MEMORY_FENCE()`** *macro. However, for portability to non-EBC environments and for readability, the use of the* **`MEMORY_FENCE()`** *macro is strongly encouraged.*

Many driver/bus master pairs do not exhibit this issue, based on the style of their driver/bus master interactions.

The easy/safe case is when the driver builds a structure and the bus master will not access that structure until the (exclusive to this request, or only one request outstanding) bus master "doorbell" is rung, and the doorbell resides on a PCI device. This mechanism works because **`PciIo->IoWrite()`** and **`PciIo->MemWrite()`** are also memory fence operations.

Another "safe" variation is a memory FIFO of host requests, and the host writes the current FIFO producer index to the bus master's doorbell. This mechanism is safe because the bus master will not attempt to read the FIFO until the corresponding index has been written, and the bus master write is the memory fence.

NOTE

This mechanism is really a variation of the "easy/safe case" above (two paragraphs before this note), because the producer index makes the doorbell write exclusive to a single request; i.e., the bus master does not "read ahead."

Another variation in the "unique doorbell per request" category is the bus master doorbell that is really a bus-master-based FIFO. The driver typically writes the address of the request to this FIFO register. This method is commonly used in the "I_2O" model but is used by several vendors without being I_2O compliant.

18.12 Helpful Tools

To catch possible issues with assigning or comparing values of different sizes, EFI drivers should always be compiled with fairly high warning levels. For example, the Microsoft Visual Studio tool chain supports the /WX and /W3 or /W4 compiler flags. The /WX flag will cause any compile time warnings to generate an error, so the build will stop when a warning is generated. The /W3 and /W4 flags set the warning level to 3 and 4 respectively. At these warning levels, any size mismatches in assignments and comparisons will generate a warning. With the /WX flag, the compile will stop when these size mismatches are detected.

If an EFI driver is being developed for an IA-32 system and is planned to be ported to the Itanium processor, then it is always a good idea to compile the EFI driver with an Itanium compiler during the development process to make sure the code is clean when validation on the Itanium processor is begun. By using the /WX and /W3 or /W4 compiler flags, any size mismatches that are generated by only 64-bit code will be detected.

19
EFI Byte Code Porting Considerations

There are a few considerations to keep in mind when writing drivers that may be ported to EBC. This chapter describes these considerations in detail and, where applicable, provides solutions to address them. If EFI drivers are implemented with these considerations in mind, then porting a native driver to EBC may simply require a recompile using the Intel® C Compiler for EFI Byte Code.

19.1 No EBC Assembly Support

The only tools that are available for EBC are a C compiler and a PE/COFF linker. There are no EBC assemblers available, and there are no plans to produce an EBC assembler. This lack of an EBC assembler is actually by design, because the EBC instruction set was optimized with a C compiler in mind. If the driver is being ported to EBC, all assembly language for IA-32 or Itanium processors must be converted to C.

19.2 No Floating Point Support

There is no floating-point support in the EBC Virtual Machine, which means that the type **`float`** is not supported in the Intel C Compiler for EFI Byte Code. If an EFI driver is being ported to EBC and requires floating-point math, then the driver must be converted to fixed-point math using integer operands and operators. At this time, no fixed-point math libraries have been ported to EBC.

19.3 No C++ Support

The Intel C Compiler for EFI Byte Code does not support C++. If there is any C++ code in an EFI driver being ported to EBC, then that C++ code must be converted to C.

19.4 EFI Data Type Sizes

The most notable difference between native and EBC code execution is that, in some cases, **`sizeof()`** is computed at runtime for EBC code, whereas **`sizeof()`** can be computed at compile time for native code. Because pointers, the EFI data types **`INTN`** and **`UINTN`**, and the C type **`long`** are different sizes on IA-32 and Itanium processors, an EBC executable must adapt to the platform type on which it is executing. Example 19-1 below shows several examples of simple and complex data types.

```
typedef enum {Red, Green, Blue} COLOR_TYPE;

typedef struct {
  UINT64  ValueU64;
  UINT32  ValueU32;
  UINT16  ValueU16;
```

```
  UINT8   ValueU8;
} FIXED_STRUCTURE;

typedef struct {
  UINTN   ValueUN;
  VOID    *Pointer;
  UINT64  ValueU64;
  UINT32  ValueU32;
} VARIABLE_STRUCTURE;

Size = sizeof (UINT64);             //  8 bytes on IA-32,  8 bytes on Itanium
Size = sizeof (UINT32);             //  4 bytes on IA-32,  4 bytes on Itanium
Size = sizeof (UINT16);             //  2 bytes on IA-32,  2 bytes on Itanium
Size = sizeof (UINT8);              //  1 byte on IA-32,  1 byte on Itanium
Size = sizeof (UINTN);              //  4 bytes on IA-32,  8 bytes on Itanium
Size = sizeof (INTN);               //  4 bytes on IA-32,  8 bytes on Itanium
Size = sizeof (COLOR_TYPE);         //  4 bytes on IA-32,  8 bytes on Itanium
Size = sizeof (VOD *);              //  4 bytes on IA-32,  8 bytes on Itanium
Size = sizeof (FIXED_STRUCTURE);    // 15 bytes on IA-32, 15 bytes on Itanium
Size = sizeof (VARIABLE_STRUCTURE); // 20 bytes on IA-32, 28 bytes on Itanium
```

Example 19-1. Size of EBC Data Types

For the types that return different sizes for IA-32 and Itanium processors, the EBC compiler generates code that computes the correct values at runtime.

19.5 CASE Statements

Because pointers and the data types **INTN** and **UINTN** are different sizes on IA-32 and Itanium processors, there is only one place where the **sizeof()** function cannot be used, which is in a **case** statement. The **sizeof()** function cannot be used in a **case** statement because the **sizeof()** function cannot be evaluated to a constant by the EBC compiler at compile time. EFI status codes such as **EFI_SUCESS** and **EFI_UNSUPPORTED** are defined using the **sizeof()** function. As a result, these values cannot be used in **case** expressions. Example 19-2 shows examples of case statements.

```
UINTN Value;

switch (Value) {
case 0:                 // Works
  break;
case sizeof (UINT16):   // Works because sizeof (UINT16) is always 2
  break;
case sizeof (UINTN):    // Compiler error because sizeof (UINTN) is not constant
  break;
case EFI_UNSUPPORTED:   // Compiler error because sizeof (UINTN) is not constant
  break;
}
```

Example 19-2. Case Statements

19.6 Stronger Type Checking

The EBC compiler performs stronger type checking than some other IA-32 and Itanium compilers. As a result, code that compiles without any errors or warnings on an IA-32 or Itanium compiler may generate warnings with the EBC compiler. Example 19-3 shows two common examples, using **gBS->AllocatePool()** and **gBS->OpenProtocol()**, from EFI drivers that will generate warnings with the EBC compiler and how these examples can be fixed.

```
typedef struct {
  UINT8   Value8;
  UINT64  Value64;
} MY_STRUCTURE;

EFI_STATUS                    Status;
EFI_DRIVER_BINDING_PROTOCOL   *This;
EFI_HANDLE                    ControllerHandle;
EFI_GUID                      gEfiBlockIoProtocolGuid;
EFI_BLOCK_IO_PROTOCOL         *BlockIo;
MY_STRUCTURE                  *MyStructure;

Status = gBS->OpenProtocol (
                ControllerHandle,
                &gEfiBlockIoProtocolGuid,
                &BlockIo,                     // Compiler warning
                This->DriverBindingHandle,
                ControllerHandle,
                EFI_OPEN_PROTOCOL_BY_DRIVER
                );

Status = gBS->OpenProtocol (
                ControllerHandle,
                &gEfiBlockIoProtocolGuid,
                (VOID **)&BlockIo,            // No compiler warning
                This->DriverBindingHandle,
                ControllerHandle,
                EFI_OPEN_PROTOCOL_BY_DRIVER
                );

Status = gBS->AllocatePool (
                EfiBootServicesData,
                sizeof (MY_STRUCTURE),
                &MyStructure                  // Compiler warning
                );

Status = gBS->AllocatePool (
                EfiBootServicesData,
                sizeof (MY_STRUCTURE),
                (VOID **)&MyStructure         // No compiler warning
                );
```

Example 19-3. Stronger Type Checking

19.7 EFI Driver Entry Point

The entry point to every EBC driver is a function called **EfiStart()**. The **EfiStart()** function performs the runtime initialization of the EFI driver and then calls **EfiMain()**. The entry point that is declared in the **Make.inf** file is always renamed to **EfiMain()**. These details are hidden from the developer by the build environment. The symbols **EfiStart()** and **EfiMain()** are used by the EBC build environment. As a result, these two reserved symbols cannot be used as function names or variable names in an EFI driver implementation.

19.8 Memory Ordering

The EBC interpreter ensures that all memory transactions are strongly ordered. EBC drivers are not required to use the **MEMORY_FENCE()** macro when strong ordering is required. However, to guarantee that an EFI driver is portable to non-EBC execution environments, the use of the **MEMORY_FENCE()** macro is strongly encouraged.

19.9 Performance Considerations

All EBC executables require an EBC Virtual Machine interpreter to be executed. Because all EBC executables are running through an interpreter, they will run slower than native EFI executables. As a result, an EFI driver that is compiled with an EBC compiler should be optimized for performance to improve the usability of the EFI driver. Chapter 14 covers speed optimization techniques that can be used to improve the performance of all EFI drivers.

20
Building EFI Drivers

This section describes how to write, compile, and package EFI drivers for the *EFI 1.10 Sample Implementation* environment.

20.1 Writing EFI Drivers

New EFI drivers are added to the EFI source tree in the **\Efi1.1\Edk\Drivers** directory. It is recommended that all EFI drivers be placed below this directory, but it is not strictly required. Additional subdirectories can be created under this directory to help organize the EFI drivers into natural groups.

To add a new EFI driver to the build environment, do the following:

- Create a new subdirectory.
- Place a **Make.inf** file along with the **.c** and **.h** files in that subdirectory.

No assembly language files are allowed in this directory. If assembly language files are required, they should be placed in processor-specific subdirectories.

Example 20-1 shows the files that are present in an EFI driver that consumes the Block I/O Protocol and produces the Disk I/O Protocol. These files are placed in a driver directory called **DiskIo**.

```
Efi1.1\
    Edk\
        Drivers\
            DiskIo\
                ComponentName.c
                DiskIo.c
                DiskIo.h
                Make.inf
```

Example 20-1. Disk I/O Driver Files

The example in Example 20-1 does not contain any processor-specific files. This absence means that this driver is designed to be portable between IA-32, Itanium architecture, and EBC. If an EFI driver requires processor-specific components, then those components can be added in subdirectories below the EFI driver's directory. Table 20-1 lists the three directory names that are reserved for the processor-specific files.

Table 20-1. Directory Names Reserved for Processor-Specific Files

Directory Name	Notes
Ia32	May contain **.c**, **.h**, and **.asm** files.
Ipf	May contain **.c**, **.h**, and **.s** files.
Ebc	May contain **.c** and **.h** files.

Example 20-2 shows the same EFI driver from Example 20-1, but it includes processor-specific files for all three supported processor types. The files **FastTransfer.asm** and **FastTransfer.s** contain optimized assembly code to improve the performance of the disk I/O driver. Because the EBC compiler does not support assembly language, the same functions that **FastTransfer.asm** and **FastTransfer.s** provide must be implemented in **FastTransfer.c** for EBC. Doing so makes the driver work on all three supported processor types, but the EFI driver takes longer to develop and is more difficult to maintain if any changes are required in the processor-specific components. If possible, an EFI driver should be implemented in C with no processor-specific files, which will reduce the development time, reduce maintenance costs, and increase portability.

```
Efi1.1\
    Edk\
        Drivers\
            DiskIo\
                ComponentName.c
                DiskIo.c
                DiskIo.h
                Make.inf
                Ia32\
                    FastTransfer.asm
                Ipf\
                    FastTransfer.s
                Ebc\
                    FastTransfer.c
```

Example 20-2. Disk I/O Driver with Processor-Specific Files

20.1.1 Make.inf File

The **Make.inf** file specifies the following:

- A list of source files
- The path to the include directories
- The path to the library directories
- The entry point for the EFI driver
- The name of the executable EFI driver image

It is not legal to use relative path names with "**..**" in any of the sections. When relative path names are used, it makes it very difficult to move a component to a different location in the source tree. Likewise, **#include** statements in **.c** and **.h** files should not use relative path names with "**..**" for the same reason.

Example 20-3 shows the contents of the **Make.inf** file for the **DiskIo** driver from Example 20-1. Example 20-4 shows the **Make.inf** file for the **DiskIo** driver from Example 20-2. The source files for the disk I/O driver are included in Appendix D for reference.

```
#
# Copyright (c)  1999 - 2003 Intel Corporation. All rights reserved
# This software and associated documentation (if any) is furnished
# under a license and may only be used or copied in accordance
# with the terms of the license. Except as permitted by such
# license, no part of this software or documentation may be
# reproduced, stored in a retrieval system, or transmitted in any
# form or by any means without the express written consent of
```

```
# Intel Corporation.
#
# Module Name:
#
#    make.inf
#
# Abstract:
#
#    Makefile for Edk\Drivers\DiskIo
#
# Revision History
#

[sources]
  DiskIo.c
  DiskIo.h
  ComponentName.c

[includes]
  .
  $(EFI_SOURCE)\Edk
  $(EFI_SOURCE)\Edk\Include
  $(EFI_SOURCE)\Edk\Lib\Include

[libraries]
  $(EFI_SOURCE)\Edk\Lib\EfiDriverLib
  $(EFI_SOURCE)\Edk\Lib\Print
  $(EFI_SOURCE)\Edk\Protocol

[nmake]
  IMAGE_ENTRY_POINT=DiskIoDriverEntryPoint
  TARGET_BS_DRIVER=DiskIo
```

Example 20-3. Disk I/O Driver Make.inf File

```
#
# Copyright (c)  1999 - 2003 Intel Corporation. All rights reserved
# This software and associated documentation (if any) is furnished
# under a license and may only be used or copied in accordance
# with the terms of the license. Except as permitted by such
# license, no part of this software or documentation may be
# reproduced, stored in a retrieval system, or transmitted in any
# form or by any means without the express written consent of
# Intel Corporation.
#
# Module Name:
#
#    make.inf
#
# Abstract:
#
#    Makefile for Edk\Drivers\DiskIo
#
# Revision History
#

[sources]
  DiskIo.c
```

```
  DiskIo.h
  ComponentName.c

[ia32sources]
  FastTransfer.asm

[ipfsources]
  FastTransfer.s

[ebcsources]
  FastTransfer.c

[includes]
  .
  $(EFI_SOURCE)\Edk
  $(EFI_SOURCE)\Edk\Include
  $(EFI_SOURCE)\Edk\Lib\Include

[libraries]
  $(EFI_SOURCE)\Edk\Lib\EfiDriverLib
  $(EFI_SOURCE)\Edk\Lib\Print
  $(EFI_SOURCE)\Edk\Protocol

[nmake]
  IMAGE_ENTRY_POINT=DiskIoDriverEntryPoint
  TARGET_BS_DRIVER=DiskIo
```

Example 20-4. Disk I/O Driver Make.inf File with Processor-Specific Files

The following sections describe the different sections that are available in the **`Make.inf`** file.

20.1.2 [sources] Section

Four sources sections are available in a **`make.inf`** file. Table 20-2 describes each of these four types of sources sections.

Table 20-2. Sources Sections Available in a make.inf File

Type of Section	Optional or Required?	Description
[sources]	Required	Contains the names of the **`.c`** and **`.h`** files that are in the EFI driver's directory. If possible, a driver should be designed so that this section is the only sources section. If any of the other sources sections are used, then additional work will be required to port the EFI driver to the other supported processor types.
[ia32sources]	Optional	Contains the names of the IA-32 specific **`.c`**, **`.h`**, and **`.asm`** files that are in the **`Ia32`** subdirectory below the EFI driver's directory.
[ipfsources]	Optional	Contains the names of the Itanium architecture–specific **`.c`**, **`.h`**, and **`.s`** files that are in the **`Ipf`** subdirectory below the EFI driver's directory.
[ebcsources]	Optional	Contains the names of the EFI Byte Code–specific **`.c`** and **`.h`** files that are in the **`Ebc`** subdirectory below the EFI driver's directory.

The example in Example 20-1 does not contain any processor-specific sources, so it contains only a single **[sources]** section. The **Make.inf** file from Example 20-2 would contain all four sources sections.

20.1.3 [includes] Section

The **[includes]** section in a **Make.inf** file contains the list of include directories that are required by the EFI driver. All EFI drivers require the include directory named "**.**" to include the **.h** files that are listed from the EFI driver's directory and the processor-specific subdirectories. All EFI drivers also require the **$(EFI_SOURCE)\Edk\Include** directory to include all the function prototypes and data structure definitions for EFI Boot Services and EFI Runtime Services. If an EFI driver uses the EFI Driver Library, then the include directory **$(EFI_SOURCES)\Edk\Lib\Include** is also required. If the EFI driver includes any protocols or GUIDs (see Appendix A), then the **$(EFI_SOURCE)\Edk** directory is also required. Additional include paths can be added as required. The example in Example 20-3 uses all four of these include directories.

20.1.4 [libraries] Section

The **[libraries]** section in a **Make.inf** file contains the list of libraries against which the EFI driver will be linked. Table 20-3 lists the libraries that are required when the EFI driver has different characteristics. If any new libraries are added to the build environment, then they can be added here too.

Table 20-3. Required Libraries

If the EFI driver…	Required Library
Uses the EFI Driver Library	**$(EFI_SOURCE)\Edk\Lib\EfiDriverLib**
Includes any protocols (see Appendix A)	**$(EFI_SOURCE)\Edk\Protocol**
Includes any GUIDs (see Appendix A)	**$(EFI_SOURCE)\Edk\Guid**
Contains any **DEBUG()** macros	**$(EFI_SOURCE)\Edk\Lib\Print**
Requires the use of print functions	**$(EFI_SOURCE)\Edk\Lib\Print**
Requires any string functions from the EFI Driver Library (see Appendix A)	**$(EFI_SOURCE)\Edk\Lib\String**

20.1.5 [nmake] Section

The **[nmake]** section in a **Make.inf** file contains two required lines, which are listed in Table 20-4. Any other statements in the **[nmake]** section are added to the **makefile** that is used to build the EFI driver. This feature can be used to override the default compiler and linker flags.

Table 20-4. Required Lines in an [nmake] Section

Required Line	Description
IMAGE_ENTRY_POINT=	Declares the EFI driver's entry point.
Either **TARGET_BS_DRIVER=** or **TARGET_RT_DRIVER=**	• The **TARGET_BS_DRIVER=** statement declares that the EFI driver is an EFI Boot Services driver. • The **TARGET_RT_DRIVER=** statement declares that the EFI driver is an EFI runtime driver.

20.2 Adding an EFI Driver to a Build Tip

Before a new EFI driver can be compiled, the **Makefile** for one or more build tips needs to be updated to include the new EFI driver. All of the build tips produce two directories when they are built. Table 20-5 lists these directories, and Example 20-5 below shows the directory structure for all the build tips.

Table 20-5. Directory Structure for All Build Tips

Directory	Description
Bin	Contains the executable images that are generated when the build tip is linked. A newly compiled EFI driver will be placed in this directory. The tree can be cleaned by deleting the **Bin** directory.
Output	Contains all the **.Obj** and **.Lib** files that are created when the build tip is compiled. This is done to prevent the **.Obj** and **.Lib** files from being scattered throughout the source tree. To clean the tree of these temporary files, the **Output** directory can be deleted.

```
Efi1.1\
    Build\
        Bios32\
            Bin\
            Output\
        EbcDrivers\
            Bin\
            Output\
        IA-32Emb\
            Bin\
            Output\
        Ia32Drivers\
            Bin\
            Output\
        IpfDrivers\
            Bin\
            Output\
```

```
        Nt32\
            Bin\
            Output\
        Sal64\
            Bin\
            Output\
```

Example 20-5. EFI Build Tips

Table 20-6 lists the standard build tips that are the fully integrated EFI builds. The build tips include the EFI core that produces the EFI Boot Services and EFI Runtime Services as well as a number of EFI drivers that provide access to a wide variety of boot devices.

Table 20-6. Build Tips Integrated into EFI Builds

Build Tips	Description
Bios32	Produces floppy images that can be booted on a PC-AT class IA-32 platform.
IA-32Emb	Produces floppy images that can be booted on a PC-AT class IA-32 platform.
Sal64	Produces the EFI firmware component for an Itanium-based platform that contains a SAL and a legacy IA-32 BIOS.
Nt32	Produces a 32-bit Windows application that is an accurate emulation of the EFI execution environment.

Table 20-7 lists the build tips that are used to build EFI drivers that are either going to be soft loaded or integrated into an option ROM container on a device. These build tips will be used during initial EFI driver development work.

Table 20-7. Build Tips Used to Build EFI Drivers

Build Tips	Description
Ia32Drivers	Produces drivers that can be executed in the **Bios32**, **IA-32Emb**, or **Nt32** environments.
IpfDrivers	Produces drivers that can be executed in the **Sal64** environment.
EbcDrivers	Produces drivers that can be executed in all four of the EFI environments that are listed in Table 20-6.

It is also possible to create additional build tips below the **\Efi1.1\Build** directory and use the other build tips as a template for setting up the build environment.

Once a build tip is selected, the **Makefile** in that build tip must be modified. Typically, there is a section called **Makemaker::** or **Drivers::** in the **Makefile**. To add a driver to the build tip, just add a line to this section with the path to the new driver. Example 20-6 shows the **Makemaker::** section before and after the XYZ EFI driver was added.

```
makemaker::
  $(MAKE) -f output\Edk\drivers\AtapiPassThru\makefile                    all
  $(MAKE) -f output\Edk\drivers\Console\ConPlatform\makefile              all
  $(MAKE) -f output\Edk\drivers\Console\ConSplitter\makefile              all
  $(MAKE) -f output\Edk\drivers\Console\GraphicsConsole\makefile          all
  $(MAKE) -f output\Edk\drivers\Console\Terminal\makefile                 all
  $(MAKE) -f Output\Edk\Drivers\CirrusLogic5430\makefile                  all
```

```
makemaker::
  $(MAKE) -f output\Edk\drivers\AtapiPassThru\makefile                    all
  $(MAKE) -f output\Edk\drivers\Console\ConPlatform\makefile              all
  $(MAKE) -f output\Edk\drivers\Console\ConSplitter\makefile              all
  $(MAKE) -f output\Edk\drivers\Console\GraphicsConsole\makefile          all
  $(MAKE) -f output\Edk\drivers\Console\Terminal\makefile                 all
  $(MAKE) -f Output\Edk\Drivers\CirrusLogic5430\makefile                  all
  $(MAKE) -f Output\Edk\Drivers\Xyz\makefile                              all
```

Example 20-6. Adding an EFI Driver to a Makefile

Once an EFI driver has been added to the **Makefile**, the EFI driver can be built by executing the **Build.cmd** and **nmake** commands from the command line. See the release notes from the *EFI 1.10 Sample Implementation* for details on the environment variables that are set up in the **Build.cmd** file. If the build completes, then the EFI driver executable will be placed in the **Bin** directory below the build tip. The name and type for the EFI driver image is specified in the **Make.inf** file (see section 20.1).

20.3 Integrating an EFI Driver into a Build Tip

If an EFI driver is being added to the **Bios32**, **IA-32Emb**, **Sal64**, or the **Nt32** environments, then the EFI driver also needs to be linked into the environment and the driver needs to be loaded and started prior to the execution of the EFI boot manager. This integration requires a few more files to be modified to integrate a new EFI driver. There are also some additional EFI driver design considerations to make sure the integration is successful. The following steps are required to integrate an EFI driver into one of the four build tips.

1. Add the EFI driver to the build tip. See section 20.2 for details on updating the **Makemaker::** section. The **EFI_LIBS** section also needs to be updated for the **Bios32**, **IA-32Emb**, and **Sa64** build tips to directly link the EFI driver into the final executable image. Example 20-7 below shows an example with the XYZ driver being added to both the **EFI_LIBS** section and the **Makemaker::** section. This step needs to be repeated for all the build tips in which the EFI driver is being integrated.

```
EFI_LIBS =                                                              \
  Output\Corefw\Fw\Efi\Efi.lib                                          \
  . . .                                                                 \
  Output\Edk\Drivers\Xyz\Xyz.lib                                        \

makemaker::
  $(MAKE) -f output\Edk\drivers\AtapiPassThru\makefile                  all
  . . .
  $(MAKE) -f Output\Edk\Drivers\Xyz\makefile                            all
```

Example 20-7. Integrating an EFI Driver to a Makefile

2. Add the declaration of the EFI driver's entry point to the file **\Efi1.1\CoreFw\Fw\Platform\BuildTip\Inc\Drivers.h**. Example 20-8 below shows a portion of the **Drivers.h** file that has been updated to include the declaration of the XYZ driver's entry point at the end of the file.

```
EFI_STATUS
InitializeEbcDriver (
  IN EFI_HANDLE       ImageHandle,
  IN EFI_SYSTEM_TABLE *SystemTable
  );

EFI_STATUS
DecompressDriverEntryPoint (
  IN EFI_HANDLE       ImageHandle,
  IN EFI_SYSTEM_TABLE *SystemTable
  );

. . .

EFI_STATUS
XyzDriverEntryPoint (
  IN EFI_HANDLE       ImageHandle,
  IN EFI_SYSTEM_TABLE *SystemTable
  );
```

Example 20-8. Adding an EFI Driver's Entry Point to Drivers.h

3. Add a call to the EFI driver's entry point to the build tip by updating the file **\Efi1.1\CoreFw\Fw\Platform\BuildTip\XXX\Init.c**, where **XXX** is either **Bios32**, **IA-32Emb**, **Sal64**, or **Nt32**. The call to the EFI driver's entry point should be added to the function called **MainEntry()**. The safest place to add this call is just before the loop that calls the EFI boot manager. However, if the driver that is being added is required to establish a console or it provides a set of services that are consumed by EFI drivers that are executed later, the call to the driver's initialization function will have to be moved closer to the beginning of **MainEntry()**. Example 20-9 below shows the XYZ driver being called just before the loop calls the EFI boot manager.

```
    //
    // Install all built in EFI 1.1 Drivers that require
    // EFI Variable Services
    //
    LOAD_INTERNAL_BS_DRIVER (L"Terminal",    InitializeTerminal);
    LOAD_INTERNAL_BS_DRIVER (L"ConPlatform", ConPlatformDriverEntry);
    LOAD_INTERNAL_BS_DRIVER (L"ConSplitter", ConSplitterDriverEntry);
    LOAD_INTERNAL_BS_DRIVER (L"Snp3264",     InitializeSnpNiiDriver);
    LOAD_INTERNAL_BS_DRIVER (L"PxeBc",       InitializeBCDriver);
    LOAD_INTERNAL_BS_DRIVER (L"BIS",         EFIBIS_BaseCodeModuleInit);
    LOAD_INTERNAL_BS_DRIVER (L"IsaFloppy",   FdcControllerDriverEntryPoint);
    LOAD_INTERNAL_BS_DRIVER (L"Ps2Mouse",    PS2MouseDriverEntryPoint);
    LOAD_INTERNAL_BS_DRIVER (L"UsbMouse",    USBMouseDriverBindingEntryPoint);
    LOAD_INTERNAL_BS_DRIVER (L"SerialMouse", SerialMouseDriverEntryPoint);

    LOAD_INTERNAL_BS_DRIVER (L"XyzDriver",   XyzDriverEntryPoint);  // Added

    //
    // Create an event to be signalled when ExitBootServices occurs
```

```
    //
    Status = BS->CreateEvent(
                    EVT_SIGNAL_EXIT_BOOT_SERVICES,
                    TPL_NOTIFY,
                    PlExitBootServices,
                    NULL,
                    &Event
                    );
    ASSERT (!EFI_ERROR(Status));

#ifdef EFI_BOOTSHELL
    PlInitializeInternalLoad();
#endif

    //
    // Loop through boot manager and boot maintenance until a boot
    // option is selected
    //
    while (TRUE) {
        //
        // The platform code is ready to boot the machine. Pass control
        // to the boot manager.
        //
        LOAD_INTERNAL_DRIVER(
            FW,
            IMAGE_SUBSYSTEM_EFI_APPLICATION,
            L"bootmgr",
            InitializeBootManager
            );

        //
        // If we return from above, it means that no boot choices were found
        // or boot maintenance was chosen. Hence invoke boot maintenance menu.
        //
        LOAD_INTERNAL_DRIVER(
            FW,
            IMAGE_SUBSYSTEM_EFI_APPLICATION,
            L"bmaint",
            InitializeBootMaintenance
            );
    }
}
```

Example 20-9. Calling an EFI Driver's Entry Point

4. There is one important factor that must be considered when integrating an EFI driver into the **Bios32**, **IA-32Emb**, or **Sal64** build tips. This consideration is to make sure there are no global function or global variable symbol collisions between the EFI driver that is being added and the EFI core or EFI drivers that are already present in the build tip. When the final link is performed, no warnings or errors will be generated because most linkers support having the same symbol in multiple libraries. This linker behavior is expected and allows newer libraries to override older libraries. However, with EFI drivers, the results can be disastrous because two different EFI drivers may use the same global variable or the same global function. To detect this issue, a line can be added to the **Makefile** for the build tip that is being integrated. This line will link all the **EFI_LIBS** together to create a merged library. This link will fail if there multiple global functions or global variables are defined. The line to add is shown in Example 20-10 below. It links all the **EFI_LIBS** together into the file **AllLibs.Lib** before the final link is performed. If a symbol that is defined multiple times is listed for the EFI driver that is being integrated, then some of the global variables or global functions must either be declared **static** or renamed.

```
$(OUTPUTS): $(EFI_LIBS)
  $(LIB) $(EFI_LIBS) /out:AllLibs.Lib
  $(LINK) $(L_FLAGS) $(NT_LIBS) $(EFI_LIBS) /entry:MainEntry /out:$@ /pdb:$*.pdb
```

Example 20-10. Linking All EFI_LIBS together in Makefile

20.4 Build Tools

This section describes the build tools that are contained in the *EFI Sample Implementation*. Some of these tools are automatically used in the build process, and some are stand-alone tools that can be used to post-process the EFI images that are generated by the build process. Table 20-8 describes the tools that are required in the build process, and Table 20-9 describes the stand-alone tools.

Table 20-8. EFI Build Tools Required by the Build Process

Tool	Description
FwImage	A utility that adjusts the subsystem field of a PE/COFF header to mark a file as an EFI application, EFI Boot Service driver, or EFI runtime driver. The build environment uses this tool internally.
GenHelp	A utility that converts a Unicode text file into a **.c** file containing an array of **CHAR16** values. This tool is used to convert the EFI Shell help information file in **\Efi1.1\Shell\HelpData\HelpData.Src** into the **HelpData.c** file during the build process.
GenMake	A utility that converts a **Make.Inf** file into a **Makefile** for each component in a build tip. The build environment uses this tool internally.
EfiLdrImage	A utility that is used to build a bootable floppy image for the **Bios32** and **IA-32Emb** build tips. The build environment uses this tool internally.
SplitFile	A utility that is used to build a bootable floppy image for the **Bios32** and **IA-32Emb** build tips. The build environment uses this tool internally.

Table 20-9. Stand-alone EFI Build Tools

Tool	Description
Col	A utility that converts TAB characters to four spaces in `.c` and `.h` source files.
DskImage	A utility that transfers a disk image file to a floppy disk. This tool is used only in the `Bios32` and `IA-32Emb` build environments.
EfiCompress	A utility that compresses a file using the EFI 1.10 compression algorithm. A file compressed using this utility can be decompressed using the services of the EFI Decompress Protocol or the EFI Shell command `EfiDecompress`. This tool is used in the `Bios32` and `IA-32Emb` build environments, but it can also be used to measure the effectiveness of the EFI 1.10 compression algorithm on EFI driver images.
EfiRom	A utility that creates a binary file that is a PCI option ROM image that conforms to the *PCI 2.2 Specification.* It takes as input raw binary images, EFI images, or any combination of the two. By default all EFI images are compressed, and the images are placed in the PCI option ROM in the order in which they are passed to this tool. The EFI Shell command `LoadPciRom` can be used to load compressed and uncompressed EFI drivers from the output files that this tool generates. The binary file that this tool generates is also the same image that can be directly programmed into a PCI option ROM.
Futil	A utility that reprograms the FLASH on an Intel® network interface controller (NIC).

20.4.1 FwImage Build Tool

EFI drivers are standard PE/COFF images with an EFI-specific subsystem type. See section 2.1.1 of the *EFI 1.10 Specification* for details on EFI image subsystem types. The build utility `Fwimage.exe` is provided in the *EFI 1.10 Sample Implementation* to change the subsystem type of a PE/COFF image to one of the supported EFI-specific subsystem types. This utility will be required until commercial tool chains build in support for the EFI-specific subsystem types. At this time, the EBC linker is the only linker that contains support for these subsystem types. The following three examples show how the `FwImage` build tool can be used to convert `foo.exe` to `foo.efi` with the subsystem types for an EFI application, EFI Boot Service driver, and EFI runtime driver.

```
fwimage app   foo.exe foo.efi
fwimage bsdrv foo.exe foo.efi
fwimage rtdrv foo.exe foo.efi
```

The following `#define` statements are the subsystem type values for EFI applications, EFI Boot Service drivers, and EFI runtime drivers from the *EFI 1.10 Sample Implementation*.

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

20.4.2 EfiRom Build Tool

The **EfiRom** build tool helps in the development of EFI drivers for PCI adapters. Once an EFI driver for a PCI adapter has been added to a build tip, it needs to be packaged into a PCI option ROM. EFI drivers stored in PCI option ROMs are automatically loaded and executed by the PCI bus driver. The PCI driver model that is documented in chapter 12 of the *EFI 1.10 Specification* provides the ability to mix legacy BIOS images and EFI driver images for IA-32, Itanium architecture, and EBC. The **EfiRom** build tool allows any combination of these image types to be combined into a single PCI option ROM image. The final output of the **EfiRom** tool can be used with a PROM programmer or a flash update utility to reprogram the option ROM device on a PCI adapter. There are many options to the **EfiRom** tool, which are listed below in Table 20-10 for reference. Below the references are several examples showing how the **EfiRom** tool can be used to build many different types of PCI option ROM images.

Table 20-10. EFIROM Tool Switches

Switches	Description
-p	Provides verbose output from this tool.
-l	Do not automatically set the LAST bit in the PCIR data structure of the last image in the output file.
-v VendorId	A required switch that specifies the 16-bit vendor ID in the PCIR data structure for all images specified with the **-e** and **-ec** switches.
-d DeviceId	A required switch that specifies the 16-bit device ID in the PCIR data structure for all images specified with the **-e** and **-ec** switches.
-cc ClassCode	An optional switch that specifies the 24-bit class code in the PCIR data structure for all images specified with the **-e** and **-ec** switches.
-rev Revision	An optional switch that specifies the 16-bit revision in the PCIR data structure for all images specified with the **-e** and **-ec** switches.
-o OutFileName	An optional switch.
-e Filenames	Specifies the list of EFI driver images that are to be added to the PCI option ROM image.
-ec Filenames	Specifies the list of EFI driver images that are to be compressed and added to the PCI option ROM image.
-b Filenames	Specifies the list of binary images that are to be added to the PCI option ROM image. If the binary file is not the last image, then the LAST bit in the PCIR data structure will be cleared. If it is the last image, then the LAST bit in the PCIR data structure will be set. If the -l switch is specified, then the LAST bit will not be modified.
-dump FileName	Dumps the contents of a PCI option ROM image that follows the *PCI 2.2 Specification* for PCI option ROM construction.

Example 20-11 shows the following examples using the **EfiRom** tool:

- **Example 1:** Builds a PCI option ROM that contains a single uncompressed EFI driver. The vendor ID is **0xABCD**, and the device ID is **0x1234**.
- **Example 2:** Same as example 1, but the EFI driver is compressed, which reduces the size of the PCI option ROM.
- **Example 3:** Same as example 2, but an output file is specified. Examples 1 and 2 will generate the output file **EfiDriver.Rom**. Example 2 generates the output file **PciOptionRom.Rom**.
- **Example 4:** Puts two compressed EFI drivers in a PCI option ROM.
- **Example 5:** Puts a binary file and a compressed EFI driver in a PCI option ROM. This example is how a legacy option ROM image can be combined with an EFI driver to produce an adapter that works on PC-AT class systems and EFI systems.
- **Example 6:** Shows how the class code and code revision fields can also be specified. By default, these fields are cleared by the **EfiRom** tool.
- **Example 7:** Shows how the PCI option ROM image generated in example 3 can be dumped.

```
(1) Efirom -v 0xabcd -d 0x1234 -e EfiDriver.efi

(2) Efirom -v 0xabcd -d 0x1234 -ec EfiDriver.efi

(3) Efirom -v 0xabcd -d 0x1234 -ec EfiDriver.efi -o PciOptionRom.Rom

(4) Efirom -v 0xabcd -d 0x1234 -ec EfiDriverA.efi EfiDriverB.efi

(5) Efirom -v 0xabcd -d 0x1234 -b Bios.bin -ec EfiDriver

(6) Efirom -v 0xabcd -d 0x1234 -cc 0x030000 -rev 0x0001 -b
Bios.bin -ec EfiDriver

(7) Efirom -dump PciOptionRom.Rom
```

Example 20-11. EFIROM Tool Examples

21
Testing and Debugging EFI Drivers

Always provide a driver in both native-instruction-set and EBC binary forms. Providing both of these forms allows the OEM firmware to simulate testing the driver in a fast best-case scenario and a slower scenario. If the driver is tested to work as both an EBC and native-instruction-set binary, it is expected that there will be fewer timing sensitivities to the driver and it will be more robust.

There are several EFI Shell commands that can be used to help debug EFI drivers. These EFI Shell commands are fully documented in the *EFI 1.1 Shell Commands Specification,* so the full capabilities of the EFI Shell commands will not be discussed here. The EFI Shell that is included in the *EFI Sample Implementation* is a reference implementation of an EFI Shell that may be customized for various platforms. As a result, the EFI Shell commands described here may not behave identically on all platforms. There is also a built-in EFI Shell command call **`help`** that will provide a detailed description of an EFI Shell command. The EFI Shell commands that are listed in Table 21-1 can be used to test and debug EFI drivers. The protocols and services exercised by each of these commands are also listed in Table 21-1.

Table 21-1. EFI Shell Commands

Command	Protocol Tested	Service Tested
Load -nc		EFI driver entry point.
Load		EFI driver entry point.
	Driver Binding	Supported()
	Driver Binding	Start()
Unload	Loaded Image	Unload()
Connect	Driver Binding	Supported()
		Start()
Disconnect	Driver Binding	Stop()
Reconnect	Driver Binding	Supported()
	Driver Binding	Start()
	Driver Binding	Stop()

continued

Table 21-1. EFI Shell Commands (continued)

Command	Protocol Tested	Service Tested
Drivers	Component Name	GetDriverName()
Devices	Component Name Component Name	GetDriverName() GetControllerName()
DevTree	Component Name	GetControllerName()
Dh -d	Component Name Component Name	GetDriverName() GetControllerName()
OpenInfo		
DrvCfg -s	Driver Configuration	SetOptions()
DrvCfg -f	Driver Configuration	ForceDefaults()
DrvCfg -v	Driver Configuration	OptionsValid()
DrvDiag	Driver Diagnostics	RunDiagnostics()
Err		

21.1 Loading EFI Drivers

Table 21-2 lists the two EFI Shell commands that are available to load and start EFI drivers.

Table 21-2. EFI Shell Commands for Loading EFI Drivers

Command	Description

Load	Loads an EFI driver from a file. EFI driver files typically have an extension of `.efi`. The `Load` command has one important option, the `-nc` ("No Connect") option, for EFI driver developers. When the `Load` command is used without the `-nc` option, then the loaded driver will automatically be connected to any devices in the system that it is able to manage. This means that the EFI driver's entry point is executed and then the EFI Boot Service `ConnectController()` is called. If the EFI driver produces the Driver Binding Protocol in the driver's entry point, then the `ConnectController()` call will exercise the `Supported()` and `Start()` services of Driver Binding Protocol that was produced.
	If the `-nc` option is used with the `Load` command, then this automatic connect operation will not be performed. Instead, only the EFI driver's entry point is executed. When the `-nc` option is used, the EFI Shell command `Connect` can be used to connect the EFI driver to any devices in the system that it is able to manage. The `Load` command can also take wild cards, so multiple EFI drivers can be loaded at the same time.
	The code below shows the following examples of the `Load` command:
	• **Example 1:** Loads and does not connect the EFI driver image `EfiDriver.efi`. This example exercises only the EFI driver's entry point.
	• **Example 2:** Loads and connects the EFI driver image called `EfiDriver.efi`. This example exercises the EFI driver's entry point and the `Supported()` and `Start()` functions of the Driver Binding Protocol.
	• **Example 3:** Loads and connects all the EFI drivers with an `.efi` extension from `fs0:`, exercising the EFI driver entry points and their `Supported()` and `Start()` functions of the Driver Binding Protocol.
	`fs0:> Load -nc EfiDriver.efi`
	`fs0:> Load EfiDriver.efi`
	`fs0:> Load *.efi`

LoadPciRom	Simulates the load of a PCI option ROM by the PCI bus driver. This command parses a ROM image that was produced with the `EfiRom` build utility. Details on the `EfiRom` build utility can be found in section 20.4.2. The `LoadPciRom` command will find all the EFI drivers in the ROM image and will attempt to load and start all the EFI drivers. This command helps test the ROM image before it is burned into a PCI adapter's ROM. No automatic connects are performed by this command, so only the EFI driver's entry point will be exercised by this command. The EFI Shell command `Connect` will have to be used for the loaded EFI drivers to start managing devices. The example below loads and calls the entry point of all the EFI drivers in the ROM file called `MyAdapter.ROM`. `fs0:> LoadPciRom MyAdapter.ROM`

21.2 Unloading EFI Drivers

Table 21-3 lists the EFI Shell commands that can be used to unload an EFI driver if it is unloadable.

Table 21-3. EFI Shell Commands for Unloading EFI Drivers

Command	Description
Unload	Unloads an EFI driver if it is unloadable. This command takes a single argument that is the image handle number of the EFI driver to unload. The `Dh -p Image` command and the `Drivers` command can be used to search for the image handle of the driver to unload. Once the image handle number is known, an unload operation can be attempted. The `Unload` command may fail for one of the following two reasons: • The EFI driver may not be unloadable, because EFI drivers are not required to be unloadable. • The EFI driver might be unloadable, but it may not be able to be unloaded right now. Some EFI drivers may need to be disconnected before they are unloaded. They can be disconnected with the `Disconnect` command. The following example unloads the EFI driver on handle 27. If the EFI driver on handle 27 is unloadable, it will have registered an `Unload()` function in its Loaded Image Protocol. This command will exercise the EFI driver's `Unload()` function. `Shell> Unload 27`

21.3 Connecting EFI Drivers

Table 21-4 lists the three EFI Shell commands that can be used to test the connecting of EFI drivers to devices. There are many options for using these commands. A few are shown in Table 21-4.

Table 21-4. EFI Shell Commands for Connecting EFI Drivers

Command	Description

Connect	Can be used to connect all EFI drivers to all devices in the system or connect EFI drivers to a single device. The code below shows the following examples of the `Connect` command:
	• **Example 1:** Connects all drivers to all devices.
	• **Example 2:** Connects all drivers to the device that is abstracted by handle 23.
	• **Example 3:** Connects the EFI driver on handle 27 to the device that is abstracted by handle 23.
	`fs0:> Connect -r`
	`fs0:> Connect 23`
	`fs0:> Connect 23 27`

continued

Table 21-4. EFI Shell Commands for Connecting EFI Drivers (continued)

Command	Description

Disconnect	Stops EFI drivers from managing a device. The code below shows the following examples of the **Disconnect** command:
	• **Example 1:** Disconnects all drivers from all devices. However, the use of this command is not recommended, because it will also disconnect all the console devices.
	• **Example 2:** Disconnects all the EFI drivers from the device represented by handle 23.
	• **Example 3:** Disconnects the EFI driver on handle 27 from the device represented by handle 23.
	• **Example 1:** Destroys the child represented by handle 32. The EFI driver on handle 27 produced the child when it started managing the device on handle 23.
	`fs0:> Disconnect -r`
	`fs0:> Disconnect 23`
	`fs0:> Disconnect 23 27`
	`fs0:> Disconnect 23 27 32`

Reconnect	Is the equivalent of executing the **Disconnect** and **Connect** commands back to back. The **Reconnect** command is the best command for testing the Driver Binding Protocol of EFI drivers. This command tests the **Supported()**, **Start()**, and **Stop()** services of the Driver Binding Protocol. The **Reconnect -r** command tests the Driver Binding Protocol for every EFI driver that follows the EFI Driver Model. Use this command before an EFI driver is loaded to verify that the current set of drivers pass the **Reconnect -r** test, and then load the new EFI driver and rerun the **Reconnect -r** test. An EFI driver is not complete until it passes this interoperability test with the EFI core and the full set of EFI drivers. The code below shows the following examples of the **Reconnect** command:
	• **Example 1:** Reconnects all the EFI drivers to the device handle 23.
	• **Example 2:** Reconnects the EFI driver on handle 27 to the device on handle 23.
	• **Example 3:** Reconnects all the EFI drivers in the system.
	`fs0:> Reconnect 23`
	`fs0:> Reconnect 23 27`
	`fs0:> Reconnect -r`

21.4 Driver and Device Information

Table 21-5 lists the EFI Shell commands that can be used to dump information about the EFI drivers that follow the EFI Driver Model. Each of these commands shows information from a slightly different perspective.

Table 21-5. EFI Shell Commands for Driver and Device Information

Command	Description

Drivers	Lists all the EFI drivers that follow the EFI Driver Model. It uses the **`GetDriverName()`** service of the Component Name protocol to retrieve the human-readable name of each EFI driver if it is available. It also shows the file path from which the EFI driver was loaded. As EFI drivers are loaded with the **`Load`** command, they will appear in the list of drivers produced by the **`Drivers`** command. The **`Drivers`** command can also show the name of the EFI driver in different languages. The code below shows the following examples of the **`Drivers`** command:
	• **Example 1:** Shows the **`Drivers`** command being used to list the EFI drivers in the default language.
	• **Example 2:** Shows the driver names in Spanish.
	`fs0:> Drivers`
	`fs0:> Drivers -lspa`

Devices	Lists all the devices that are being managed or produced by EFI drivers that follow the EFI Driver Model. This command uses the **`GetControllerName()`** service of the Component Name protocol to retrieve the human-readable name of each device that is being managed or produced by EFI drivers. If a human-readable name is not available, then the EFI device path is used. The code below shows the following examples of the **`Devices`** command:
	• **Example 1:** Shows the **`Devices`** command being used to list the EFI drivers in the default language.
	• **Example 2:** Shows the device names in Spanish.
	`fs0:> Drivers`
	`fs0:> Drivers -lspa`

continued

Table 21-5. EFI Shell Commands for Driver and Device Information (continued)

Command	Description

DevTree	Similar to the **Devices** command. Lists all the devices being managed by EFI drivers that follow the EFI Driver Model. This command uses the **GetControllerName()** service of the Component Name Protocol to retrieve the human-readable name of each device that is being managed or produced by EFI drivers. If the human-readable name is not available, then the EFI device path is used. This command also visually shows the parent/child relationships between all of the devices by displaying them in a tree structure. The lower a device is in the tree of devices, the more the device name is indented. The code below shows the following examples of the **DevTree** command:
	• **Example 1:** Displays the device tree with the device names in the default language.
	• **Example 2:** Displays the device tree with the device names in Spanish.
	• **Example 3:** Displays the device tree with the device names shown as EFI device paths.
	`fs0:> DevTree`
	`fs0:> DevTree -lspa`
	`fs0:> DevTree -d`

Dh -d	Provides a more detailed view of a single driver or a single device than the **Drivers**, **Devices**, and **DevTree** commands. If a driver binding handle is used with the **Dh -d** command, then a detailed description of that EFI driver is provided along with the devices that the driver is managing and the child devices that the driver has produced. If a device handle is used with the **Dh -d** command, then a detailed description of that device is provided along with the drivers that are managing that device, that device's parent controllers, and the device's child controllers. If the **Dh -d** command is used without any parameters, then detailed information on all of the drivers and devices is displayed. The code below shows the following examples of the **Dh -d** command:
	• **Example 1:** Displays the details on the EFI driver on handle 27.
	• **Example 2:** Displays the details for the device on handle 23.
	• **Example 3:** Shows details on all the drivers and devices in the system.
	`fs0:> Dh -d 27`
	`fs0:> Dh -d 23`
	`fs0:> Dh -d`

OpenInfo	Provides detailed information on a device handle that is being managed by one or more EFI drivers that follow the EFI Driver Model. The **OpenInfo** command displays each protocol interface installed on the device handle, and the list of agents that have opened that protocol interface with the **OpenProtocol()** Boot Service. This command can be used in conjunction with the **Connect**, **Disconnect**, and **Reconnect** commands to verify that an EFI driver is opening and closing protocol interfaces correctly. The following example shows the **OpenInfo** command being used to display the list of protocol interfaces on device handle 23 along with the list of agents that have opened those protocol interfaces. `fs0:> OpenInfo 23`

21.5 Testing the Driver Configuration Protocol

Table 21-6 lists the EFI Shell commands that can be used to test the Driver Configuration Protocol.

Table 21-6. EFI Shell Commands for Testing the Driver Configuration Protocol

Command	Description

DrvCfg	Provides the services that are required to test the Driver Configuration Protocol implementation of an EFI driver. This command can show all the devices that are being managed by EFI drivers that support the Driver Configuration Protocol. The **Devices** and **Drivers** commands will also show the drivers that support the Driver Configuration Protocol and the devices that those drivers are managing or have produced. Once a device has been chosen, the **DrvCfg** command can be used to invoke the **SetOptions()**, **ForceDefaults()**, or **OptionsValid()** services of the Driver Configuration Protocol. The code below shows the following examples of the **DrvCfg** command:
	• **Example 1:** Displays all the devices that are being managed by EFI drivers that support the Driver Configuration Protocol.
	• **Example 2:** Forces defaults on all the devices in the system.
	• **Example 3:** Validates the options on all the devices in the system.
	• **Example 4:** Invokes the **SetOptions()** service of the Driver Configuration Protocol for the driver on handle 27 and the device on handle 23.
	`fs0:> DrvCfg`
	`fs0:> DrvCfg -f`
	`fs0:> DrvCfg -v`
	`fs0:> DrvCfg -s 23 27`

21.6 Testing the Driver Diagnostics Protocol

Table 21-7 lists the EFI Shell commands that can be used to test the Driver Diagnostics Protocol.

Table 21-7. EFI Shell Commands for Testing the Driver Diagnostics Protocol

Command	Description

DrvDiag	Provides the ability to test all the services of the Driver Diagnostics Protocol that is produced by an EFI driver. This command is able to show the devices that are being managed by EFI drivers that support the Driver Diagnostics Protocol. The **`Devices`** and **`Drivers`** commands will also show the drivers that support the Driver Diagnostics Protocol and the devices that those drivers are managing or have produced. Once a device has been chosen, the **`DrvDiag`** command can be used to invoke the **`RunDiagnostics()`** service of the Driver Diagnostics Protocol. The code below shows the following examples of the **`DrvDiag`** command:
	• **Example 1:** Displays all the devices that are being managed by EFI drivers that support the Driver Diagnostics Protocol.
	• **Example 2:** Invokes the **`RunDiagnostics()`** service of the Driver Diagnostics Protocol in standard mode for the driver on handle 27 and the device on handle 23.
	• **Example 3:** Invokes the **`RunDiagnostics()`** service of the Driver Diagnostics Protocol in manufacturing mode for the driver on handle 27 and the device on handle 23.
	`fs0:> DrvDiag`
	`fs0:> DrvDiag -s 23 27`
	`fs0:> DrvDiag -m 23 27`

21.7 ASSERT() and DEBUG() Macros

Every module will have a debug (check) build and a clean build. The debug build will include code for debug that will not be included in normal clean production builds. A debug build is enabled when the identifier **EFI_DEBUG** exists. A clean build is defined as when the **EFI_DEBUG** identifier does not exist.

The following debug macros can be used to insert debug code into a checked build. This debug code can greatly reduce the amount of time it takes to root cause a bug. These macros are enabled only in a debug build, so they do not take up any executable space in the production build. Table 21-8 describes the debug macros that are available.

Table 21-8. Available Debug Macros

Debug Macro	Description
ASSERT (Expression)	For check builds, if **Expression** evaluates to **FALSE**, a diagnostic message is printed and the program is aborted. Aborting a program is usually done via the **EFI_BREAKPOINT ()** macro. For clean builds, **Expression** does not exist in the program and no action is taken. Code that is required for normal program execution should never be placed inside an **ASSERT** macro, as the code will not exist in a production build.
ASSERT_EFI_ERROR (Status)	For check builds, an assert is generated if **Status** is an error. This macro is equivalent to **ASSERT (!EFI_ERROR (Status))** but is easier to read.
DEBUG ((ErrorLevel, String, …))	For check builds, **String** and its associated arguments will be printed if the **ErrorLevel** of the macro is active. See Table 21-9 below for a definition of the **ErrorLevel** values.
DEBUG_CODE (Code)	For check builds, **Code** is included in the build. **DEBUG_CODE (** is on its own line and indented like normal code. All the debug code follows on subsequent lines and is indented an extra level. The **)** is on the line following all the code and indented at the same level as **DEBUG_CODE (**.
EFI_BREAKPOINT ()	On a check build, inserts a break point into the code.
DEBUG_SET_MEM (Address, Length)	For a check build, initializes the memory starting at **Address** for **Length** bytes with the value **BAD_POINTER**. This initialization is done to enable debug of code that uses memory buffers that are not initialized.

continued

Table 21-8. Available Debug Macros (continued)

Debug Macro	Description
`CR (Record, TYPE, Field, Signature)`	The containing record macro returns a pointer to `TYPE` when given the structure's `Field` name and a pointer to it (`Record`). The `CR` macro returns the `TYPE` pointer for check and production builds. For a check build, an `ASSERT ()` is generated if the `Signature` field of `TYPE` is not equal to the `Signature` in the `CR ()` macro.

The `ErrorLevel` parameter referenced in the `DEBUG()` macro allows an EFI driver to assign a different error level to each debug message. Critical errors should always be sent to the standard error device. However, informational messages that are used only to debug a driver should be sent to the standard error device only if the user wants to see those specific types of messages. The EFI Shell supports the `Err` command that allows the user to set the error level. The EFI Boot Maintenance Manager allows the user to enable and select a standard error device. It is recommended that a serial port be used as a standard error device during debug so the messages can be logged to a file with a terminal emulator. Table 21-9 contains the list of error levels that are supported in the *EFI Sample Implementation*.

Table 21-9. Error Levels

Mnemonic	Value	Description
EFI_D_INIT	0x00000001	Initialization messages
EFI_D_WARN	0x00000002	Warning messages
EFI_D_LOAD	0x00000004	Load events
EFI_D_FS	0x00000008	EFI file system messages
EFI_D_POOL	0x00000010	EFI pool allocation and free messages
EFI_D_PAGE	0x00000020	EFI page allocation and free messages
EFI_D_INFO	0x00000040	Informational messages
EFI_D_VARIABLE	0x00000100	EFI variable service messages
EFI_D_BM	0x00000400	EFI boot manager messages
EFI_D_BLKIO	0x00001000	EFI Block I/O Protocol messages
EFI_D_NET	0x00004000	EFI Simple Network Protocol, PXE base code, BIS messages
EFI_D_UNDI	0x00010000	UNDI driver messages
EFI_D_LOADFILE	0x00020000	Load File Protocol messages
EFI_D_EVENT	0x00080000	EFI Event Services messages
EFI_D_ERROR	0x80000000	Critical error messages

21.8 POST Codes

If an EFI driver is being developed that cannot make use of the **DEBUG()** and **ASSERT()** macros, then a different mechanism must be used to help in the debugging process. Under these conditions, it is usually sufficient to send a small amount of output to a device to indicate what portions of an EFI driver have executed and where error conditions have been detected. A few possibilities are presented below, but many others are possible depending on the devices that may be available on a specific platform. It is important to note that these mechanisms are useful during driver development and debug, but they should never be present in production versions of EFI drivers because these types of devices are not present on all platforms.

There are a couple of possibilities. The first is to use a POST card. A POST card is an ISA or PCI adapter that displays the hex value of an 8-bit I/O write cycle to address 0x80. If an EFI driver can depend on the PCI Root Bridge I/O Protocol being present, then the driver can use the services of the PCI Root Bridge I/O Protocol to send an 8-bit I/O write cycle to address 0x80. A driver can also use the services of the PCI I/O Protocol to write to address 0x80, as long as the pass-through BAR value is used. Example 21-1 below shows how the PCI Root Bridge I/O and PCI I/O Protocols can be used to send a value to a POST card.

```
EFI_STATUS                        Status;
EFI_PCI_ROOT_BRIDGE_IO_PROTOCOL   *PciRootBridgeIo;
EFI_PCI_IO_PROTOCOL               *PciIo;
UINT8                             Value;

Value = 0xAA;
Status = PciRootBridgeIo->Io.Write (
                                PciRootBridgeIo,
                                EfiPciWidthUint8,
                                0x80,
                                1,
                                &Value
                                );

Value = 0xAA;
Status = PciIo->Io.Write (
                      PciIo,
                      EfiPciIoWidthUint8,
                      EFI_PCI_IO_PASS_THROUGH_BAR,
                      0x80,
                      1,
                      &Value
                      );
```

Example 21-1. POST Code Examples

If a POST card is not available, then the next possibility is a text-mode VGA frame buffer. If a system initialized the text-mode VGA display by default before the EFI driver executed, then the EFI driver can make use of the PCI Root Bridge I/O or PCI I/O Protocols to directly write text characters to the text-mode VGA display. Example 21-2 shows how the PCI Root Bridge I/O and PCI I/O Protocols can be used to send the text message "ABCD" to the text-mode VGA frame buffer. Some systems do not have a VGA controller, so this solution will not work on all systems.

```
EFI_STATUS                        Status;
EFI_PCI_ROOT_BRIDGE_IO_PROTOCOL  *PciRootBridgeIo;
EFI_PCI_IO_PROTOCOL              *PciIo;
UINT8                            *Value;

Value = {'A',0x0f,'B',0x0f,'C',0x0f,'D',0x0f};

Status = PciRootBridgeIo->Mem.Write (
                                PciRootBridgeIo,
                                EfiPciWidthUint8,
                                0xB8000,
                                8,
                                Value
                                );

Status = PciIo->Mem.Write (
                      PciIo,
                      EfiPciIoWidthUint8,
                      EFI_PCI_IO_PASS_THROUGH_BAR,
                      0xB8000,
                      8,
                      Value
                      );
```

Example 21-2. VGA Display Examples

Another option is to use some type of byte-stream-based device. This device could include a UART or an SMBus, for example. Like the POST card, the idea is to use the services of the PCI Root Bridge I/O or PCI I/O Protocols to initialize and send characters to the byte-stream device.

Appendix A
EFI Data Types

Table A-1 contains the set of base types that are used in all EFI applications and EFI drivers. These are the base types that should be used to build more complex unions and structures. The file **EFIBIND.H** in the *EFI 1.10 Sample Implementation* contains the code that is required to map compiler-specific data types to the EFI data types. If a new compiler is used, only this one file should be updated; all other EFI-related sources should compile unmodified. Table A-2 contains the modifiers that can be used in conjunction with the EFI data types.

Table A-1. Common EFI Data Types

Mnemonic	Description
BOOLEAN	Logical Boolean. 1-byte value containing a 0 for **FALSE** or a 1 for **TRUE**. Other values are undefined.
INTN	Signed value of native width. (4 bytes on IA-32, 8 bytes on Itanium architecture operations)
UINTN	Unsigned value of native width. (4 bytes on IA-32, 8 bytes on Itanium architecture operations)
INT8	1-byte signed value.
UINT8	1-byte unsigned value.
INT16	2-byte signed value.
UINT16	2-byte unsigned value.
INT32	4-byte signed value.
UINT32	4-byte unsigned value.
INT64	8-byte signed value.
UINT64	8-byte unsigned value.
CHAR8	1-byte character.
CHAR16	2-byte character. Unless otherwise specified, all strings are stored in the UTF-16 encoding format as defined by Unicode 2.1 and ISO/IEC 10646 standards.
VOID	Undeclared type.
EFI_GUID	128-bit buffer containing a unique identifier value. Unless otherwise specified, aligned on a 64-bit boundary.
EFI_STATUS	Status code. Type INTN.
EFI_HANDLE	A collection of related interfaces. Type VOID *.
EFI_EVENT	Handle to an event structure. Type VOID *.
EFI_LBA	Logical block address. Type UINT64.
EFI_TPL	Task priority level. Type UINTN.

continued

Table A-1. Common EFI Data Types (continued)

Mnemonic	Description
EFI_MAC_ADDRESS	32-byte buffer containing a network Media Access Controller address.
EFI_Ipv4_ADDRESS	4-byte buffer. An IPv4 internet protocol address.
EFI_Ipv6_ADDRESS	16-byte buffer. An IPv6 internet protocol address.
EFI_IP_ADDRESS	16-byte buffer aligned on a 4-byte boundary. An IPv4 or IPv6 internet protocol address.
<Enumerated Type>	Element of an enumeration. Type INTN.

Table A-2. Modifiers for Common EFI Data Types

Mnemonic	Description
IN	Datum is passed to the function.
OUT	Datum is returned from the function.
OPTIONAL	Datum that is passed to the function is optional, and a **NULL** may be passed if the value is not supplied.
STATIC	The function has local scope. This modifier replaces the standard C static key word, so it can be overloaded for debugging.
VOLATILE	Declares a variable to be volatile and thus exempt from optimization to remove redundant or unneeded accesses. Any variable that represents a hardware device should be declared as **VOLATILE**.
CONST	Declares a variable to be of type **const**. This modifier is a hint to the compiler to enable optimization and stronger type checking at compile time.

Appendix B
EFI Status Codes

Table B-1 contains the set of status code values that may be returned by EFI Boot Services, EFI Runtime Services, and EFI protocol services. These status codes are defined in detail in Appendix D of the *EFI 1.10 Specification*.

Table B-1. EFI_STATUS Codes

Mnemonic	IA-32	Itanium Architecture
EFI_SUCCESS	0x00000000	0x0000000000000000
EFI_LOAD_ERROR	0x80000001	0x8000000000000001
EFI_INVALID_PARAMETER	0x80000002	0x8000000000000002
EFI_UNSUPPORTED	0x80000003	0x8000000000000003
EFI_BAD_BUFFER_SIZE	0x80000004	0x8000000000000004
EFI_BUFFER_TOO_SMALL	0x80000005	0x8000000000000005
EFI_NOT_READY	0x80000006	0x8000000000000006
EFI_DEVICE_ERROR	0x80000007	0x8000000000000007
EFI_WRITE_PROTECTED	0x80000008	0x8000000000000008
EFI_OUT_OF_RESOURCES	0x80000009	0x8000000000000009
EFI_VOLUME_CORRUPTED	0x8000000A	0x800000000000000A
EFI_VOLUME_FULL	0x8000000B	0x800000000000000B
EFI_NO_MEDIA	0x8000000C	0x800000000000000C
EFI_MEDIA_CHANGED	0x8000000D	0x800000000000000D
EFI_NOT_FOUND	0x8000000E	0x800000000000000E
EFI_ACCESS_DENIED	0x8000000F	0x800000000000000F
EFI_NO_RESPONSE	0x80000010	0x8000000000000010
EFI_NO_MAPPING	0x80000011	0x8000000000000011
EFI_TIMEOUT	0x80000012	0x8000000000000012
EFI_NOT_STARTED	0x80000013	0x8000000000000013
EFI_ALREADY_STARTED	0x80000014	0x8000000000000014
EFI_ABORTED	0x80000015	0x8000000000000015
EFI_ICMP_ERROR	0x80000016	0x8000000000000016

continued

Table B-1. EFI_STATUS Codes (continued)

Mnemonic	IA-32	Itanium Architecture
EFI_TFTP_ERROR	0x80000017	0x8000000000000017
EFI_PROTOCOL_ERROR	0x80000018	0x8000000000000018
EFI_INCOMPATIBLE_VERSION	0x80000019	0x8000000000000019
EFI_SECURITY_VIOLATION	0x8000001A	0x800000000000001A
EFI_CRC_ERROR	0x8000001B	0x800000000000001B
EFI_WARN_UNKOWN_GLYPH	0x00000001	0x0000000000000001
EFI_WARN_DELETE_FAILURE	0x00000002	0x0000000000000002
EFI_WARN_WRITE_FAILURE	0x00000003	0x0000000000000003
EFI_WARN_BUFFER_TOO_SMALL	0x00000004	0x0000000000000004

Appendix C
Quick Reference Guide

This appendix contains a summary of the services and GUIDs that are available to EFI drivers, including the following:

- EFI Boot Services
- EFI Runtime Services
- EFI Driver Library Services
- GUID variables
- Protocol variables
- Various protocol services

Some of the GUIDs and protocols listed here are not part of the *EFI 1.10 Specification*. Instead, they are extensions that are included in the *EFI Sample Implementation*.

The EFI Boot Services and EFI Runtime Services are listed in Table C-1, Table C-2, and Table C-3. These three tables list the most commonly used services, the rarely used services, and the services that should not be used from EFI drivers. Services labeled with type **BS** are EFI Boot Services, and services labeled with type **RT** are EFI Runtime Services. A detailed explanation of the services that are rarely used or should not be used by EFI drivers is included in chapter 2. Chapters 5 and 6 of the *EFI 1.10 Specification* contain detailed descriptions of all the EFI Boot Services and EFI Runtime Services.

Table C-1. EFI Services That Are Commonly Used by EFI Drivers

Type	Service	Type	Service
BS	gBS->AllocatePool()	BS	gBS->InstallMultipleProtocolInterfaces()
BS	gBS->FreePool()	BS	gBS->UninstallMultipleProtocolInterfaces()
BS	gBS->AllocatePages()	BS	gBS->LocateHandleBuffer()
BS	gBS->FreePages()	BS	gBS->LocateProtocol()
BS	gBS->SetMem()	BS	gBS->OpenProtocol()
BS	gBS->CopyMem()	BS	gBS->CloseProtocol()
		BS	gBS->OpenProtocolInformation()
BS	gBS->CreateEvent()		
BS	gBS->CloseEvent()	BS	gBS->RaiseTPL()
BS	gBS->SignalEvent()	BS	gBS->RestoreTPL()
BS	gBS->SetTimer()		
BS	gBS->CheckEvent()	BS	gBS->Stall()

Table C-2. EFI Services That Are Rarely Used by EFI Drivers

Type	Service	Type	Service
BS	gBS->ReinstallProtocolInterface()	BS	gBS->LoadImage()
BS	gBS->LocateDevicePath()	BS	gBS->StartImage()
		BS	gBS->UnloadImage()
BS	gBS->ConnectController()	BS	gBS->Exit()
BS	gBS->DisconnectController()		
		BS	gBS->InstallConfigurationTable()
RT	gRT->GetVariable()		
RT	gRT->SetVariable()	RT	gRT->GetTime()
BS	gBS->GetNextMonotonicCount()	BS	gBS->CalculateCrc32()
RT	gRT->GetNextHighMonotonicCount()		
		RT	gRT->ConvertPointer()

Table C-3. EFI Services That Should Not Be Used by EFI Drivers

Type	Service	Type	Service
BS	gBS->GetMemoryMap()	RT	gRT->SetVirtualAddressMap()
BS	gBS->ExitBootServices()	RT	gRT->GetNextVariableName()
BS	gBS->InstallProtocolInterface()	RT	gRT->SetTime()
BS	gBS->UninstallProtocolInterface()	RT	gRT->GetWakeupTime()
BS	gBS->HandleProtocol()	RT	gRT->SetWakeupTime()
BS	gBS->LocateHandle()		
BS	gBS->RegisterProtocolNotify()	RT	gRT->ResetSystem()
BS	gBS->ProtocolsPerHandle()	BS	gBS->SetWatchDogTimer()
BS	gBS->WaitForEvent()		

Table C-4 lists the EFI Driver Library Services. These services simplify the implementation of EFI drivers. Most of these services are layered on top of the EFI Boot Services and EFI Runtime Services. A detailed description of these library services can be found in the *EFI Driver Library Specification*.

Table C-4. EFI Driver Library Functions

Initialization Functions	**Link List Functions**	**Memory Functions**
EfiInitializeDriverLib()	InitializeListHead()	EfiCopyMem()
EfiLibInstallDriverBinding()	IsListEmpty()	EfiSetMem()
EfiLibInstallAllDriverProtocols()	RemoveEntryList()	EfiZeroMem()
	InsertTailList()	EfiCompareMem()
Device Path Functions	InsertHeadList()	EfiLibAllocatePool()
EfiDevicePathInstance()	SwapListEntries()	EfiLibAllocateZeroPool()
EfiAppendDevicePath()		
EfiAppendDevicePathNode()	**Math Functions**	**String Functions**
EfiAppendDevicePathInstance()	LShiftU64()	EfiStrCpy()
EfiFileDevicePath()	RShiftU64()	EfiStrLen()
EfiDevicePathSize()	MultU64x32()	EfiStrSize()
EfiDuplicateDevicePath()	DivU64x32()	EfiStrCmp()
		EfiStrCat()
Miscellaneous Functions	**Spin Lock Functions**	EfiLibLookupUnicodeString()
EfiCompareGuid()	EfiInitializeLock()	EfiLibAddUnicodeString()
EfiLibCreateProtocolNotifyEvent()	EfiAcquireLock()	EfiLibFreeUnicodeStringTable()
EfiLibGetSystemConfigurationTable()	EfiAcquireLockOrFail()	
	EfiReleaseLock()	

Table C-5 lists the macros that are available to EFI drivers, and Table C-6 lists the constants that are available to EFI drivers. These macros and constants simplify the implementation of EFI drivers.

Table C-5. EFI Macros

Miscellaneous Macros	Linked List Macros	Critical Error Macros
EFI_DRIVER_ENTRY_POINT ()	INITIALIZE_LIST_HEAD ()	ASSERT ()
CR ()	IS_LIST_EMPTY ()	ASSERT_EFI_ERROR ()
MEMORY_FENCE ()	REMOVE_ENTRY_LIST ()	EFI_BREAKPOINT ()
EFI_ERROR ()	INSERT_TAIL_LIST ()	ASSERT_LOCKED ()
EFI_FIELD_OFFSET ()	INSERT_HEAD_LIST ()	EFI_DEADLOOP ()
EFI_SIGNATURE_16 ()	SWAP_LIST_ENTRIES ()	
EFI_SIGNATURE_32 ()		**Debug Macros**
EFI_SIGNATURE_64 ()	**Memory Macros**	DEBUG ()
	EFI_SIZE_TO_PAGES ()	DEBUG_CODE ()
Math Macros	EFI_PAGES_TO_SIZE ()	DEBUG_SET_MEM ()
EFI_MIN ()	ALIGN_POINTER ()	
EFI_MAX ()		

Table C-6. EFI Constants

Mnemonic	Description
TRUE	One.
FALSE	Zero.
NULL	**VOID** pointer to zero.

Table C-7 lists the GUIDs that are available to EFI drivers in the *EFI Sample Implementation*. The **Directory Name** heading is the name of the GUID directory that can be used with the **EFI_GUID_DEFINITION()** macro. The **GUID Variable Names** heading contains the names of the global variables that are available to the EFI driver that uses the **EFI_GUID_DEFINITION()** macro for that directory. For example, assume the following statement is added to an EFI driver.

```
#include EFI_GUID_DEFINITION(PcAnsi)
```

Then, the following variables of type **EFI_GUID** will be available to that EFI driver:

```
gEfiPcAnsiGuid  gEfiVT100Guid  gEfiVT100PlusGuid  gEfiVTUTF8Guid
```

Table C-7. EFI GUID Variables

Directory Name	GUID Variable Names
Bmp	`gEfiDefaultBmpLogoGuid`
ConsoleInDevice	`gEfiConsoleInDeviceGuid`
ConsoleOutDevice	`gEfiConsoleOutDeviceGuid`
DebugImageInfoTable	`gEfiDebugImageInfoTableGuid`
GlobalVariable	`gEfiGlobalVariableGuid`
Gpt	`gEfiPartTypeUnusedGuid`
	`gEfiPartTypeSystemPartGuid`
	`gEfiPartTypeLegacyMbrGuid`
PcAnsi	`gEfiPcAnsiGuid`
	`gEfiVT100Guid`
	`gEfiVT100PlusGuid`
	`gEfiVTUTF8Guid`
PciOptionRomTable	`gEfiPciOptionRomTableGuid`
PrimaryConsoleInDevice	`gEfiPrimaryConsoleInDeviceGuid`
PrimaryConsoleOutDevice	`gEfiPrimaryConsoleOutDeviceGuid`
PrimaryStandardErrorDevice	`gEfiPrimaryStandardErrorDeviceGuid`
SalSystemTable	`gEfiSalSystemTableGuid`
SmBios	`gEfiSmbiosTableGuid`
StandardErrorDevice	`gEfiStandardErrorDeviceGuid`

Table C-8 lists the protocols that are available to EFI drivers in the *EFI Sample Implementation*. The **Directory Name** heading is the name of the protocol directory that can be used with the `EFI_PROTOCOL_DEFINITION()` macro. The **Protocol GUID Variable Name** heading contains the names of the global variables that are available to a driver that uses the `EFI_PROTOCOL_DEFINITION()` macro for that directory. For example, assume the following statement is added to an EFI driver.

```
#include EFI_PROTOCOL_DEFINITION(PciIo)
```

Then, the following variable of type `EFI_GUID`, along with the definition of the `EFI_PCI_IO_PROTOCOL`, would be available to that EFI driver:

```
gEfiPciIoProtocolGuid
```

Table C-8. EFI Protocol Variables

Directory Name	Protocol GUID Variable Names
BIS	gEfiBisProtocolGuid
BlockIo	gEfiBlockIoProtocolGuid
BusSpecificDriverOverride	gEfiBusSpecificDriverOverrideProtocolGuid
ComponentName	gEfiComponentNameProtocolGuid
DebugPort	gEfiDebugPortProtocolGuid
DebugSupport	gEfiDebugSupportProtocolGuid
Decompress	gEfiDecompressProtocolGuid
DeviceIo	gEfiDeviceIoProtocolGuid
DevicePath	gEfiDevicePathProtocolGuid
DiskIo	gEfiDiskIoProtocolGuid
DriverBinding	gEfiDriverBindingProtocolGuid
DriverConfiguration	gEfiDriverConfigurationProtocolGuid
DriverDiagnostics	gEfiDriverDiagnosticsProtocolGuid
Ebc	GEfiEbcProtocolGuid
	gEfiEbcDebugHelperProtocolGuid
EfiNetworkInterfaceIdentifier	gEfiNetworkInterfaceIdentifierProtocolGuid
IsaAcpi	gEfiIsaAcpiProtocolGuid
IsaIo	gEfiIsaIoProtocolGuid
LegacyBoot	gEfiLegacyBootProtocolGuid
LoadedImage	gEfiLoadedImageProtocolGuid
LoadFile	gEfiLoadFileProtocolGuid
PciIo	GEfiPciIoProtocolGuid
PciRootBridgeIo	gEfiPciRootBridgeIoProtocolGuid
PlatformDriverOverride	gEfiPlatformDriverOverrideProtocolGuid
PxeBaseCode	gEfiPxeBaseCodeProtocolGuid
PxeBaseCodeCallBack	gEfiPxeBaseCodeCallbackProtocolGuid
ScsiPassThru	gEfiScsiPassThruProtocolGuid
SerialIo	gEfiSerialIoProtocolGuid
SimpleFileSystem	gEfiSimpleFileSystemProtocolGuid
	gEfiFileInfoGuid
	GEfiFileInfoIdGuid
	GEfiFileSystemVolumeLabelInfoIdGuid

continued

Table C-8. EFI Protocol Variables (continued)

Directory Name	Protocol GUID Variable Names
SimpleNetwork	`GEfiSimpleNetworkProtocolGuid`
SimplePointer	`gEfiSimplePointerProtocolGuid`
SimpleTextIn	`gEfiSimpleTextInProtocolGuid`
SimpleTextOut	`gEfiSimpleTextOutProtocolGuid`
UgaDraw	`gEfiUgaDrawProtocolGuid`
UgaIo	`gEfiUgaIoProtocolGuid`
UgaSplash	`gEfiUgaSplashProtocolGuid`
UnicodeCollation	`gEfiUnicodeCollationProtocolGuid`
UsbAtapi	`gEfiUsbAtapiProtocolGuid`
UsbHostController	`gEfiUsbHcProtocolGuid`
UsbIo	`gEfiUsbIoProtocolGuid`
VgaMiniPort	`gEfiVgaMiniPortProtocolGuid`
WinNtIo	`gEfiWinNtIoProtocolGuid`
	`gEfiWinNtVirtualDisksGuid`
	`gEfiWinNtPhysicalDisksGuid`
	`gEfiWinNtFileSystemGuid`
	`gEfiWinNtSerialPortGuid`
	`gEfiWinNtUgaGuid`
	`gEfiWinNtConsoleGuid`
WinNtThunk	`gefiWinNtThunkProtocolGuid`

Table C-9 contains the list of the protocols that are available to EFI drivers along with the list of services that each protocol provides. This table does not show any of the data elements that a protocol interface may contain. The *EFI 1.10 Specification* should be referenced for additional details on each protocol. Each protocol is named by its C data structure. These protocols are available to EFI drivers that use the **`EFI_PROTOCOL_DEFINITION()`** macro for a specific protocol. For example, if an EFI driver intends to consume or produce the EFI USB I/O Protocol, it would need to include the following statement in its implementation:

```
#include EFI_PROTOCOL_DEFINITION(UsbIo)
```

Table C-9. EFI Protocol Service Summary

EFI_BIS_PROTOCOL	**EFI_DEVICE_IO_PROTOCOL**
Initialize()	Mem.Read()
Shutdown()	Mem.Write()
Free()	Io.Read()
GetBootObjectAuthorizationCertificate()	Io.Write()
GetBootObjectAuthorizationCheckFlag()	Pci.Read()
GetBootObjectAuthorizationUpdateToken()	Pci.Write()
GetSignatureInfo()	PciDevicePath()
UpdateBootObjectAuthorization()	Map()
VerifyBootObject()	Unmap()
VerifyObjectWithCredential()	AllocateBuffer()
EFI_BLOCK_IO_PROTOCOL	Flush()
Reset()	FreeBuffer()
ReadBlocks()	**EFI_DEVICE_PATH_PROTOCOL**
WriteBlocks()	**EFI_DISK_IO_PROTOCOL**
FlushBlocks()	ReadDisk()
EFI_BUS_SPECIFIC_DRIVER_OVERRIDE_PROTOCOL	WriteDisk()
GetDriver()	**EFI_DRIVER_BINDING_PROTOCOL**
EFI_COMPONENT_NAME_PROTOCOL	Supported()
GetDriverName()	Start()
GetControllerName()	Stop()
EFI_DEBUGPORT_PROTOCOL	**EFI_DRIVER_CONFIGURATION_PROTOCOL**
Reset()	SetOptions()
Write()	OptionsValid()
Read()	ForceDefaults()
Poll()	**EFI_DRIVER_DIAGNOSTICS_PROTOCOL**
EFI_DEBUG_SUPPORT_PROTOCOL	RunDiagnostics()
GetMaximumProcessorIndex()	**EFI_EBC_PROTOCOL**
RegisterPeriodicCallback()	CreateThunk()
RegisterExceptionCallback()	UnloadImage()
InvalidateInstructionCache()	RegisterICacheFlush()
EFI_DECOMPRESS_PROTOCOL	**EFI_FILE Handle**
GetInfo()	Open()
Decompress()	Close()

continued

Table C-9. EFI Protocol Service Summary (continued)

EFI_FILE Handle (continued)
Delete()
Read()
Write()
SetPosition()
GetPosition()
GetInfo()
SetInfo()
Flush()
EFI_ISA_ACPI_PROTOCOL
DeviceEnumerate()
SetPower()
GetCurResource()
GetPosResource()
SetResource()
EnableDevice()
InitDevice()
InterfaceInit()
EFI_ISA_IO_PROTOCOL
Mem.Read()
Mem.Write()
Io.Read()
Io.Write()
CopyMem()
Map()
Unmap()
AllocateBuffer()
FreeBuffer()
Flush()
EFI_LEGACY_BOOT_PROTOCOL
BootIt()
EFI_LOADED_IMAGE_PROTOCOL
Unload()
EFI_LOAD_FILE_PROTOCOL
LoadFile()

EFI_NETWORK_INTERFACE_IDENTIFIER_ PROTOCOL
EFI_PCI_IO_PROTOCOL
PollMem()
PollIo()
Mem.Read()
Mem.Write()
Io.Read()
Io.Write()
Pci.Read()
Pci.Write()
CopyMem()
Map()
Unmap()
AllocateBuffer()
FreeBuffer()
Flush()
GetLocation()
Attributes()
GetBarAttributes()
SetBarAttributes()
EFI_PCI_ROOT_BRIDGE_IO_PROTOCOL
PollMem()
PollIo()
Mem.Read()
Mem.Write()
Io.Read()
Io.Write()
Pci.Read()
Pci.Write()
CopyMem()
Map()
Unmap()
AllocateBuffer()
FreeBuffer()
Flush()

continued

Table C-9. EFI Protocol Service Summary (continued)

EFI_PCI_ROOT_BRIDGE_IO_PROTOCOL (continued)	**EFI_SCSI_PASS_THRU Protocol**
GetAttributes()	Mode()
SetAttributes()	PassThru()
Configuration()	GetNextDevice()
EFI_PLATFORM_DRIVER_OVERRIDE_PROTOCOL	BuildDevicePath()
GetDriver()	GetTargetLun()
GetDriverPath()	ResetChannel()
DriverLoaded()	ResetTarget()
EFI_PXE_BASE_CODE_PROTOCOL	**EFI_SIMPLE_FILE_SYSTEM_PROTOCOL**
Start()	OpenVolume()
Stop()	**EFI_SIMPLE_NETWORK_PROTOCOL**
Dhcp()	Start()
Discover()	Stop()
Mtftp()	Initialize()
UdpWrite()	Reset()
UdpRead()	Shutdown()
SetIpFilter()	ReceiveFilters()
Arp()	StationAddress()
SetParameters()	Statistics()
SetStationIp()	MCastIPtoMAC()
SetPackets()	NvData()
EFI_PXE_BASE_CODE_CALLBACK_PROTOCOL	GetStatus()
Callback()	Transmit()
PassThru()	Receive()
GetNextDevice()	**EFI_SIMPLE_POINTER_PROTOCOL**
BuildDevicePath()	Reset()
GetTargetLun()	GetState()
ResetChannel()	**EFI_SIMPLE_TEXT_IN_PROTOCOL**
ResetTarget()	Reset()
SERIAL_IO_PROTOCOL	ReadKeyStroke()
Reset()	**EFI_SIMPLE_TEXT_OUT_PROTOCOL**
SetAttributes()	Reset()
SetControl()	OutputString()
GetControl()	TestString()
Write()	QueryMode()
Read()	SetMode()

continued

Table C-9. EFI Protocol Service Summary (continued)

EFI_SIMPLE_TEXT_OUT_PROTOCOL (continued)	**EFI_USB_HC_PROTOCOL** (continued)
SetAttribute()	ControlTransfer()
ClearScreen()	BulkTransfer()
SetCursorPosition()	AsyncInterruptTransfer()
EnableCursor()	SyncInterruptTransfer()
EFI_UGA_DRAW_PROTOCOL	IsochronousTransfer()
GetMode()	AsyncIsochronousTransfer()
SetMode()	SetRootHubPortFeature()
Blt()	ClearRootHubPortFeature()
EFI_UGA_IO_PROTOCOL	**EFI_USB_IO Protocol**
DispatchService()	UsbControlTransfer()
CreateDevice()	UsbBulkTransfer()
DeleteDevice()	UsbAsyncInterruptTransfer()
EFI_UGA_SPLASH_PROTOCOL	UsbSyncInterruptTransfer()
UNICODE_COLLATION Protocol	UsbIsochronousTransfer()
StriColl()	UsbAsyncIsochronousTransfer()
MetaiMatch()	UsbGetDeviceDescriptor()
StrLwr()	UsbGetConfigDescriptor()
StrUpr()	UsbGetInterfaceDescriptor()
FatToStr()	UsbGetEndpointDescriptor()
StrToFat()	UsbGetStringDescriptor()
EFI_USB_ATAPI_PROTOCOL	UsbGetSupportedLanguages()
UsbAtapiPacketCmd()	UsbPortReset()
UsbAtapiReset()	**EFI_VGA_MINIPORT_PROTOCOL**
EFI_USB_HC_PROTOCOL	SetMode()
Reset()	**EFI_WIN_NT_IO_PROTOCOL**
GetState()	**EFI_WIN_NT_THUNK_PROTOCOL**
SetState()	Too many to list here.

Table C-10 contains the list of debug error level that can be used with the **DEBUG()** macros.

Table C-10.Error Levels

Mnemonic	Value	Description
EFI_D_INIT	0x00000001	Initialization messages
EFI_D_WARN	0x00000002	Warning messages
EFI_D_LOAD	0x00000004	Load events
EFI_D_FS	0x00000008	EFI File System messages
EFI_D_POOL	0x00000010	EFI pool allocation and free messages
EFI_D_PAGE	0x00000020	EFI page allocation a free messages
EFI_D_INFO	0x00000040	Informational messages
EFI_D_VARIABLE	0x00000100	EFI variable service messages
EFI_D_BM	0x00000400	EFI boot manager messages
EFI_D_BLKIO	0x00001000	EFI Block I/O Protocol messages
EFI_D_NET	0x00004000	EFI Simple Network Protocol, PXE Base Code, BIS messages
EFI_D_UNDI	0x00010000	UNDI driver messages
EFI_D_LOADFILE	0x00020000	Load File Protocol messages
EFI_D_EVENT	0x00080000	EFI Event Services messages
EFI_D_ERROR	0x80000000	Critical error messages

Appendix D
Disk I/O Protocol and Disk I/O Driver

This appendix contains the source files to the Disk I/O Protocol, the EFI Global Variable GUID, and the source files to the disk I/O driver. The disk I/O driver is a device driver that consumes the Block I/O Protocol and produces the Disk I/O Protocol on the same device handle.

The Disk I/O Protocol is composed of the two files **DiskIo.h** and **DiskIo.c**. **DiskIo.h** contains the GUID, function prototypes, and the data structure for the Disk I/O Protocol. The **DiskIo.c** file contains the global variable for the Disk I/O Protocol GUID.

The EFI Global Variable GUID is composed of the two files **GlobalVariable.h** and **GlobalVariable.c**. **GlobalVariable.h** contains the definition of the GUID that is used to access EFI environment variables with the EFI variable services, and the **GlobalVariable.c** file contains the declaration of the global variable for GUID defined in **GlobalVariable.h**.

The disk I/O driver is composed of the three files **DiskIo.h**, **DiskIo.c**, and **ComponentName.c**. **DiskIo.h** contains the include statements for the EFI Boot Services, EFI Runtime Services, the EFI Driver Library, the list of protocols that the disk I/O driver consumes, the list of protocols that the disk I/O driver produces, the private context structure, and the declarations of the driver's global variable. The private context structure contains the Disk I/O Protocol that is installed onto the device handle and a set of private data fields. The **DiskIo.c** file contains the disk I/O driver's entry point, the implementation of the Driver Binding Protocol functions, and the implementation of the Disk I/O Protocol functions. The **ComponentName.c** file contains the implementation of the Component Name Protocol functions.

D.1 Disk I/O Protocol - DiskIo.h

```
/*++

Copyright (c)  1999 - 2003 Intel Corporation. All rights reserved
This software and associated documentation (if any) is furnished
under a license and may only be used or copied in accordance
with the terms of the license. Except as permitted by such
license, no part of this software or documentation may be
reproduced, stored in a retrieval system, or transmitted in any
form or by any means without the express written consent of
Intel Corporation.

Module Name:

  DiskIo.h

Abstract:

  Disk IO protocol as defined in the EFI 1.0 specification.

  The Disk IO protocol is used to convert block oriented devices into byte
  oriented devices. The Disk IO protocol is intended to layer on top of the
  Block IO protocol.
```

```
--*/

#ifndef __DISK_IO_H__
#define __DISK_IO_H__

#define EFI_DISK_IO_PROTOCOL_GUID \
  { 0xce345171, 0xba0b, 0x11d2,  0x8e, 0x4f, 0x0, 0xa0, 0xc9, 0x69, 0x72, 0x3b }

EFI_INTERFACE_DECL (_EFI_DISK_IO_PROTOCOL);

typedef
EFI_STATUS
(EFIAPI *EFI_DISK_READ) (
  IN EFI_DISK_IO_PROTOCOL          *This,
  IN UINT32                        MediaId,
  IN UINT64                        Offset,
  IN UINTN                         BufferSize,
  OUT VOID                         *Buffer
  )
/*++

  Routine Description:
    Read BufferSize bytes from Offset into Buffer.

  Arguments:
    This       - Protocol instance pointer.
    MediaId    - Id of the media, changes every time the media is replaced.
    Offset     - The starting byte offset to read from
    BufferSize - Size of Buffer
    Buffer     - Buffer containing read data

  Returns:
    EFI_SUCCES            - The data was read correctly from the device.
    EFI_DEVICE_ERROR      - The device reported an error while performing the
                            read.
    EFI_NO_MEDIA          - There is no media in the device.
    EFI_MEDIA_CHNAGED     - The MediaId does not matched the current device.
    EFI_INVALID_PARAMETER - The read request contains device addresses that are
                            not valid for the device.

--*/
;

typedef
EFI_STATUS
(EFIAPI *EFI_DISK_WRITE) (
  IN EFI_DISK_IO_PROTOCOL          *This,
  IN UINT32                        MediaId,
  IN UINT64                        Offset,
  IN UINTN                         BufferSize,
  IN VOID                          *Buffer
  )
/*++

  Routine Description:
    Read BufferSize bytes from Offset into Buffer.

  Arguments:
```

```
    This       - Protocol instance pointer.
    MediaId    - Id of the media, changes every time the media is replaced.
    Offset     - The starting byte offset to read from
    BufferSize - Size of Buffer
    Buffer     - Buffer containing read data

  Returns:
    EFI_SUCCES            - The data was written correctly to the device.
    EFI_WRITE_PROTECTED   - The device cannot be written to.
    EFI_DEVICE_ERROR      - The device reported an error while performing the
                            write.
    EFI_NO_MEDIA          - There is no media in the device.
    EFI_MEDIA_CHNAGED     - The MediaId does not matched the current device.
    EFI_INVALID_PARAMETER - The write request contains device addresses that are
                            not valid for the device.

--*/
;

#define EFI_DISK_IO_PROTOCOL_REVISION 0x00010000

typedef struct _EFI_DISK_IO_PROTOCOL {
  UINT64          Revision;
  EFI_DISK_READ   ReadDisk;
  EFI_DISK_WRITE  WriteDisk;
} EFI_DISK_IO_PROTOCOL;

extern EFI_GUID gEfiDiskIoProtocolGuid;

#endif
```

D.2 Disk I/O Protocol - DiskIo.c

```
/*++

Copyright (c)  1999 - 2003 Intel Corporation. All rights reserved
This software and associated documentation (if any) is furnished
under a license and may only be used or copied in accordance
with the terms of the license. Except as permitted by such
license, no part of this software or documentation may be
reproduced, stored in a retrieval system, or transmitted in any
form or by any means without the express written consent of
Intel Corporation.

Module Name:

  DiskIo.c

Abstract:

  Disk IO protocol as defined in the EFI 1.0 specification.

  The Disk IO protocol is used to convert block oriented devices into byte
  oriented devices. The Disk IO protocol is intended to layer on top of the
  Block IO protocol.

--*/
```

```
#include "Efi.h"
#include EFI_PROTOCOL_DEFINITION (DiskIo)

EFI_GUID gEfiDiskIoProtocolGuid = EFI_DISK_IO_PROTOCOL_GUID;

EFI_GUID_STRING(
  &gEfiDiskIoProtocolGuid,
  "DiskIo Protocol",
  "EFI 1.0 Disk IO Protocol"
  );
```

D.3 EFI Global Variable GUID – GlobalVariable.h

```
/*++

Copyright (c)  1999 - 2003 Intel Corporation. All rights reserved
This software and associated documentation (if any) is furnished
under a license and may only be used or copied in accordance
with the terms of the license. Except as permitted by such
license, no part of this software or documentation may be
reproduced, stored in a retrieval system, or transmitted in any
form or by any means without the express written consent of
Intel Corporation.

Module Name:

    GlobalVariable.h

Abstract:

  GUID for EFI (NVRAM) Variables. Defined in EFI 1.0.

--*/

#ifndef _GLOBAL_VARIABLE_GUID_H_
#define _GLOBAL_VARIABLE_GUID_H_

#define EFI_GLOBAL_VARIABLE_GUID  \
  { 0x8BE4DF61, 0x93CA, 0x11d2, 0xAA, 0x0D, 0x00, 0xE0, 0x98, 0x03, 0x2B, 0x8C }

extern EFI_GUID gEfiGlobalVariableGuid;

#endif
```

D.4 EFI Global Variable GUID – GlobalVariable.c

```
/*++

Copyright (c)  1999 - 2003 Intel Corporation. All rights reserved
This software and associated documentation (if any) is furnished
under a license and may only be used or copied in accordance
with the terms of the license. Except as permitted by such
license, no part of this software or documentation may be
reproduced, stored in a retrieval system, or transmitted in any
form or by any means without the express written consent of
Intel Corporation.
```

```
Module Name:

    GlobalVariable.c

Abstract:

  GUID for EFI (NVRAM) Variables. Defined in EFI 1.0.

--*/

#include "Efi.h"
#include EFI_GUID_DEFINITION (GlobalVariable)

EFI_GUID gEfiGlobalVariableGuid = EFI_GLOBAL_VARIABLE_GUID;

EFI_GUID_STRING(&gEfiGlobalVariableGuid, "Efi", "Efi Variable GUID")
```

D.5 Disk I/O Driver - DiskIo.h

```
/*++

Copyright (c)  1999 - 2003 Intel Corporation. All rights reserved
This software and associated documentation (if any) is furnished
under a license and may only be used or copied in accordance
with the terms of the license. Except as permitted by such
license, no part of this software or documentation may be
reproduced, stored in a retrieval system, or transmitted in any
form or by any means without the express written consent of
Intel Corporation.


Module Name:

  DiskIo.h

Abstract:
  Private Data definition for Disk IO driver

--*/

#ifndef _DISK_IO_H
#define _DISK_IO_H

#include "Efi.h"
#include "EfiDriverLib.h"

//
// Driver Consumed Protocol Prototypes
//
#include EFI_PROTOCOL_DEFINITION (DevicePath)
#include EFI_PROTOCOL_DEFINITION (BlockIo)

//
// Driver Produced Protocol Prototypes
//
#include EFI_PROTOCOL_DEFINITION (DriverBinding)
#include EFI_PROTOCOL_DEFINITION (ComponentName)
#include EFI_PROTOCOL_DEFINITION (DiskIo)
```

```
#define DISK_IO_PRIVATE_DATA_SIGNATURE  EFI_SIGNATURE_32 ('d','s','k','I')

typedef struct {
  UINTN                   Signature;
  EFI_DISK_IO_PROTOCOL    DiskIo;
  EFI_BLOCK_IO_PROTOCOL   *BlockIo;
  UINT32                  BlockSize;
} DISK_IO_PRIVATE_DATA;

#define DISK_IO_PRIVATE_DATA_FROM_THIS(a) CR (a, DISK_IO_PRIVATE_DATA, DiskIo,
DISK_IO_PRIVATE_DATA_SIGNATURE)

//
// Global Variables
//
extern EFI_DRIVER_BINDING_PROTOCOL gDiskIoDriverBinding;
extern EFI_COMPONENT_NAME_PROTOCOL gDiskIoComponentName;

#endif
```

D.6 Disk I/O Driver - DiskIo.c

```
/*++

Copyright (c) 2000 Intel Corporation

Module Name:

  DiskIo.c

Abstract:

  DiskIo driver that layers itself on every Block IO protocol in the system.
  DiskIo converts a block oriented device to a byte oriented device.

  ReadDisk may have to do reads that are not aligned on sector boundaries.
  There are three cases:

    UnderRun - The first byte is not on a sector boundary or the read request is
               less than a sector in length.

    Aligned  - A read of N contiguous sectors.

    OverRun  - The last byte is not on a sector boundary.

--*/

#include "DiskIo.h"

//
// Prototypes
// Driver model protocol interface
//

EFI_STATUS
DiskIoDriverEntryPoint (
```

```
  IN EFI_HANDLE           ImageHandle,
  IN EFI_SYSTEM_TABLE     *SystemTable
  );

EFI_STATUS
EFIAPI
DiskIoDriverBindingSupported (
  IN EFI_DRIVER_BINDING_PROTOCOL    *This,
  IN EFI_HANDLE                     ControllerHandle,
  IN EFI_DEVICE_PATH_PROTOCOL       *RemainingDevicePath
  );

EFI_STATUS
EFIAPI
DiskIoDriverBindingStart (
  IN EFI_DRIVER_BINDING_PROTOCOL    *This,
  IN EFI_HANDLE                     ControllerHandle,
  IN EFI_DEVICE_PATH_PROTOCOL       *RemainingDevicePath
  );

EFI_STATUS
EFIAPI
DiskIoDriverBindingStop (
  IN  EFI_DRIVER_BINDING_PROTOCOL    *This,
  IN  EFI_HANDLE                     ControllerHandle,
  IN  UINTN                          NumberOfChildren,
  IN  EFI_HANDLE                     *ChildHandleBuffer
  );

//
// Disk I/O Protocol Interface
//

EFI_STATUS
EFIAPI
DiskIoReadDisk (
  IN EFI_DISK_IO_PROTOCOL  *This,
  IN UINT32                MediaId,
  IN UINT64                Offset,
  IN UINTN                 BufferSize,
  OUT VOID                 *Buffer
  );

EFI_STATUS
EFIAPI
DiskIoWriteDisk (
  IN EFI_DISK_IO_PROTOCOL  *This,
  IN UINT32                MediaId,
  IN UINT64                Offset,
  IN UINTN                 BufferSize,
  IN VOID                  *Buffer
  );

static EFI_DRIVER_BINDING_PROTOCOL mDiskIoDriverBinding = {
  DiskIoDriverBindingSupported,
  DiskIoDriverBindingStart,
  DiskIoDriverBindingStop,
  1,
  NULL,
```

```
  NULL
};

EFI_DRIVER_ENTRY_POINT (DiskIoDriverEntryPoint)

EFI_STATUS
EFIAPI
DiskIoDriverEntryPoint (
  IN EFI_HANDLE        ImageHandle,
  IN EFI_SYSTEM_TABLE  *SystemTable
  )
/*++

Routine Description:
  Register Driver Binding protocol for this driver.

Arguments:
  (Standard EFI Image entry - EFI_IMAGE_ENTRY_POINT)

Returns:
  EFI_SUCCESS - Driver loaded.
  other       - Driver not loaded.

--*/
{
  return EfiLibInstallDriverBinding (
           ImageHandle,
           SystemTable,
           &mDiskIoDriverBinding,
           ImageHandle
           );
}


EFI_STATUS
EFIAPI
DiskIoDriverBindingSupported (
  IN EFI_DRIVER_BINDING_PROTOCOL  *This,
  IN EFI_HANDLE                   ControllerHandle,
  IN EFI_DEVICE_PATH_PROTOCOL     *RemainingDevicePath  OPTIONAL
  )
/*++

  Routine Description:
    Test to see if this driver supports ControllerHandle. Any ControllerHandle
    than contains a BlockIo protocol can be supported.

  Arguments:
    This                - Protocol instance pointer.
    ControllerHandle    - Handle of device to test.
    RemainingDevicePath - Not used.

  Returns:
    EFI_SUCCESS         - This driver supports this device.
    EFI_ALREADY_STARTED - This driver is already running on this device.
    other               - This driver does not support this device.

--*/
{
```

```
  EFI_STATUS              Status;
  EFI_BLOCK_IO_PROTOCOL  *BlockIo;

  //
  // Open the IO Abstraction(s) needed to perform the supported test.
  //
  Status = gBS->OpenProtocol (
                 ControllerHandle,
                 &gEfiBlockIoProtocolGuid,
                 &BlockIo,
                 This->DriverBindingHandle,
                 ControllerHandle,
                 EFI_OPEN_PROTOCOL_BY_DRIVER
                 );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Close the I/O Abstraction(s) used to perform the supported test.
  //
  Status = gBS->CloseProtocol (
                 ControllerHandle,
                 &gEfiBlockIoProtocolGuid,
                 This->DriverBindingHandle,
                 ControllerHandle
                 );
  return Status;
}

EFI_STATUS
EFIAPI
DiskIoDriverBindingStart (
  IN EFI_DRIVER_BINDING_PROTOCOL  *This,
  IN EFI_HANDLE                   ControllerHandle,
  IN EFI_DEVICE_PATH_PROTOCOL     *RemainingDevicePath  OPTIONAL
  )
/*++

  Routine Description:
    Start this driver on ControllerHandle by opening a Block IO protocol and
    installing a Disk IO protocol on ControllerHandle.

  Arguments:
    This                - Protocol instance pointer.
    ControllerHandle    - Handle of device to bind driver to.
    RemainingDevicePath - Not used, always produce all possible children.

  Returns:
    EFI_SUCCESS         - This driver is added to ControllerHandle.
    EFI_ALREADY_STARTED - This driver is already running on ControllerHandle.
    other               - This driver does not support this device.

--*/
{
  EFI_STATUS              Status;
  EFI_BLOCK_IO_PROTOCOL   *BlockIo;
  DISK_IO_PRIVATE_DATA    *Private;
```

```
  Private = NULL;

  //
  // Connect to the Block IO interface on ControllerHandle.
  //
  Status = gBS->OpenProtocol (
                 ControllerHandle,
                 &gEfiBlockIoProtocolGuid,
                 &BlockIo,
                 This->DriverBindingHandle,
                 ControllerHandle,
                 EFI_OPEN_PROTOCOL_BY_DRIVER
                 );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Initialize the Disk IO device instance.
  //
  Status = gBS->AllocatePool(
                 EfiBootServicesData,
                 sizeof (DISK_IO_PRIVATE_DATA),
                 &Private
                 );
  if (EFI_ERROR (Status)) {
    goto ErrorExit;
  }
  EfiZeroMem (Private, sizeof(DISK_IO_PRIVATE_DATA));

  Private->Signature       = DISK_IO_PRIVATE_DATA_SIGNATURE;
  Private->BlockIo         = BlockIo;
  Private->BlockSize       = BlockIo->Media->BlockSize;

  Private->DiskIo.Revision  = EFI_DISK_IO_PROTOCOL_REVISION;
  Private->DiskIo.ReadDisk  = DiskIoReadDisk;
  Private->DiskIo.WriteDisk = DiskIoWriteDisk;

  //
  // Install protocol interfaces for the Disk IO device.
  //
  Status = gBS->InstallProtocolInterface (
                 &ControllerHandle,
                 &gEfiDiskIoProtocolGuid,
                 EFI_NATIVE_INTERFACE,
                 &Private->DiskIo
                 );
  if (!EFI_ERROR (Status)) {
    return EFI_SUCCESS;
  }

ErrorExit:
  gBS->CloseProtocol (
         ControllerHandle,
         &gEfiBlockIoProtocolGuid,
         This->DriverBindingHandle,
         ControllerHandle
         );
  return Status;
```

```
}

EFI_STATUS
EFIAPI
DiskIoDriverBindingStop (
  IN  EFI_DRIVER_BINDING_PROTOCOL    *This,
  IN  EFI_HANDLE                     ControllerHandle,
  IN  UINTN                          NumberOfChildren,
  IN  EFI_HANDLE                     *ChildHandleBuffer
  )
/*++

  Routine Description:
    Stop this driver on ControllerHandle by removing Disk IO protocol and closing
    the Block IO protocol on ControllerHandle.

  Arguments:
    This              - Protocol instance pointer.
    ControllerHandle  - Handle of device to stop driver on.
    NumberOfChildren  - Not used.
    ChildHandleBuffer - Not used.

  Returns:
    EFI_SUCCESS         - This driver is removed ControllerHandle.
    other               - This driver was not removed from this device.

--*/
{
  EFI_STATUS            Status;
  EFI_DISK_IO_PROTOCOL  *DiskIo;
  DISK_IO_PRIVATE_DATA           *Private;

  //
  // Get our context back.
  //
  Status = gBS->OpenProtocol (
                 ControllerHandle,
                 &gEfiDiskIoProtocolGuid,
                 &DiskIo,
                 This->DriverBindingHandle,
                 ControllerHandle,
                 EFI_OPEN_PROTOCOL_GET_PROTOCOL
                 );
  if (EFI_ERROR (Status)) {
    return EFI_UNSUPPORTED;
  }

  Private = DISK_IO_PRIVATE_DATA_FROM_THIS (DiskIo);

  Status = gBS->UninstallProtocolInterface (
                 ControllerHandle,
                 &gEfiDiskIoProtocolGuid,
                 &Private->DiskIo
                 );
  if (!EFI_ERROR (Status)) {

    Status = gBS->CloseProtocol (
                   ControllerHandle,
                   &gEfiBlockIoProtocolGuid,
```

```
                    This->DriverBindingHandle,
                    ControllerHandle
                    );
  }

  if (!EFI_ERROR (Status)) {
    gBS->FreePool (Private);
  }

  return Status;
}

EFI_STATUS
EFIAPI
DiskIoReadDisk (
  IN EFI_DISK_IO_PROTOCOL  *This,
  IN UINT32                MediaId,
  IN UINT64                Offset,
  IN UINTN                 BufferSize,
  OUT VOID                 *Buffer
  )
/*++

  Routine Description:
    Read BufferSize bytes from Offset into Buffer.

    Reads may support reads that are not aligned on
    sector boundaries. There are three cases:

      UnderRun - The first byte is not on a sector boundary or the read request
                 is less than a sector in length.

      Aligned  - A read of N contiguous sectors.

      OverRun  - The last byte is not on a sector boundary.


  Arguments:
    This       - Protocol instance pointer.
    MediaId    - Id of the media, changes every time the media is replaced.
    Offset     - The starting byte offset to read from.
    BufferSize - Size of Buffer.
    Buffer     - Buffer containing read data.

  Returns:
    EFI_SUCCESS           - The data was read correctly from the device.
    EFI_DEVICE_ERROR      - The device reported an error while performing the
                            read.
    EFI_NO_MEDIA          - There is no media in the device.
    EFI_MEDIA_CHNAGED     - The MediaId does not matched the current device.
    EFI_INVALID_PARAMETER - The read request contains device addresses that are
                            not valid for the device.

--*/
{
  EFI_STATUS    Status;
  DISK_IO_PRIVATE_DATA   *Private;
  UINT64        Lba;
  UINT64        OverRunLba;
```

```
UINTN         UnderRun;
UINTN         OverRun;
BOOLEAN       TransactionComplete;
UINTN         WorkingBufferSize;
UINT8         *WorkingBuffer;
UINTN         Length;
UINT8         *Data;

Private = DISK_IO_PRIVATE_DATA_FROM_THIS (This);

if (Private->BlockIo->Media->MediaId != MediaId) {
  return EFI_MEDIA_CHANGED;
}

WorkingBuffer = Buffer;
WorkingBufferSize = BufferSize;

Status = gBS->AllocatePool (
                EfiBootServicesData,
                Private->BlockSize,
                &Data
                );

if (EFI_ERROR (Status)) {
  goto Done;
}

Lba = DivU64x32 (Offset, Private->BlockSize, &UnderRun);

Length = Private->BlockSize - UnderRun;
TransactionComplete = FALSE;

Status = EFI_SUCCESS;
if (UnderRun != 0) {
  //
  // Offset starts in the middle of an Lba, so read the entire block.
  //
  Status = Private->BlockIo->ReadBlocks (
                              Private->BlockIo,
                              MediaId,
                              Lba,
                              Private->BlockSize,
                              Data
                              );

  if (EFI_ERROR (Status)) {
    goto Done;
  }

  if (Length > BufferSize) {
    Length = BufferSize;
    TransactionComplete = TRUE;
  }

  EfiCopyMem (WorkingBuffer, Data + UnderRun, Length);

  WorkingBuffer += Length;

  WorkingBufferSize  -= Length;
```

```
    if (WorkingBufferSize == 0) {
      goto Done;
    }

    Lba += 1;
  }


  OverRunLba = Lba + DivU64x32 (WorkingBufferSize, Private->BlockSize, &OverRun);

  if (!TransactionComplete && WorkingBufferSize >= Private->BlockSize) {
    //
    // If the DiskIo maps directly to a BlockIo device do the read.
    //
    if (OverRun != 0) {
      WorkingBufferSize -= OverRun;
    }

    Status = Private->BlockIo->ReadBlocks (
                                Private->BlockIo,
                                MediaId,
                                Lba,
                                WorkingBufferSize,
                                WorkingBuffer
                                );

    WorkingBuffer += WorkingBufferSize;
  }


  if (!TransactionComplete && OverRun != 0) {
    //
    // Last read is not a complete block.
    //
    Status = Private->BlockIo->ReadBlocks (
                                Private->BlockIo,
                                MediaId,
                                OverRunLba,
                                Private->BlockSize,
                                Data
                                );

    if (EFI_ERROR (Status)) {
      goto Done;
    }

    EfiCopyMem (WorkingBuffer, Data, OverRun);
  }

Done:
  if (Data) {
    gBS->FreePool (Data);
  }

  return Status;
}

EFI_STATUS
EFIAPI
```

```
DiskIoWriteDisk (
  IN EFI_DISK_IO_PROTOCOL  *This,
  IN UINT32                MediaId,
  IN UINT64                Offset,
  IN UINTN                 BufferSize,
  IN VOID                  *Buffer
  )
/*++

  Routine Description:
    Read BufferSize bytes from Offset into Buffer.

    Writes may require a read modify write to support writes that are not
    aligned on sector boundaries. There are three cases:

      UnderRun - The first byte is not on a sector boundary or the write request
                 is less than a sector in length. Read modify write is required.

      Aligned  - A write of N contiguous sectors.

      OverRun  - The last byte is not on a sector boundary. Read modified write
                 required.

  Arguments:
    This       - Protocol instance pointer.
    MediaId    - Id of the media, changes every time the media is replaced.
    Offset     - The starting byte offset to read from.
    BufferSize - Size of Buffer.
    Buffer     - Buffer containing read data.

  Returns:
    EFI_SUCCESS           - The data was written correctly to the device.
    EFI_WRITE_PROTECTED   - The device cannot be written to.
    EFI_DEVICE_ERROR      - The device reported an error while performing the
                            write.
    EFI_NO_MEDIA          - There is no media in the device.
    EFI_MEDIA_CHNAGED     - The MediaId does not matched the current device.
    EFI_INVALID_PARAMETER - The write request contains device addresses that are
                            not valid for the device.

--*/
{
  EFI_STATUS    Status;
  DISK_IO_PRIVATE_DATA   *Private;
  UINT64        Lba;
  UINT64        OverRunLba;
  UINTN         UnderRun;
  UINTN         OverRun;
  BOOLEAN       TransactionComplete;
  UINTN         WorkingBufferSize;
  UINT8         *WorkingBuffer;
  UINTN         Length;
  UINT8         *Data;

  Private = DISK_IO_PRIVATE_DATA_FROM_THIS (This);

  if (Private->BlockIo->Media->ReadOnly) {
    return EFI_WRITE_PROTECTED;
  }
```

```
if (Private->BlockIo->Media->MediaId != MediaId) {
  return EFI_MEDIA_CHANGED;
}

Status = gBS->AllocatePool (
                EfiBootServicesData,
                Private->BlockSize,
                &Data
                );

if (EFI_ERROR (Status)) {
  goto Done;
}

WorkingBuffer = Buffer;
WorkingBufferSize = BufferSize;

Lba = DivU64x32 (Offset, Private->BlockSize, &UnderRun);

Length = Private->BlockSize - UnderRun;
TransactionComplete = FALSE;

Status = EFI_SUCCESS;
if (UnderRun != 0) {
  //
  // Offset starts in the middle of an Lba, so do read modify write.
  //
  Status = Private->BlockIo->ReadBlocks (
                              Private->BlockIo,
                              MediaId,
                              Lba,
                              Private->BlockSize,
                              Data
                              );

  if (EFI_ERROR (Status)) {
    goto Done;
  }

  if (Length > BufferSize) {
    Length = BufferSize;
    TransactionComplete = TRUE;
  }

  EfiCopyMem (Data + UnderRun, WorkingBuffer, Length);

  Status = Private->BlockIo->WriteBlocks (
                              Private->BlockIo,
                              MediaId,
                              Lba,
                              Private->BlockSize,
                              Data
                              );
  if (EFI_ERROR (Status)) {
    goto Done;
  }
```

```
    WorkingBuffer += Length;
    WorkingBufferSize  -= Length;
    if (WorkingBufferSize == 0) {
      goto Done;
    }

    Lba += 1;
  }

  OverRunLba = Lba + DivU64x32 (WorkingBufferSize, Private->BlockSize, &OverRun);

  if (!TransactionComplete && WorkingBufferSize >= Private->BlockSize) {
    //
    // If the DiskIo maps directly to a BlockIo device do the write.
    //
    if (OverRun != 0) {
      WorkingBufferSize -= OverRun;
    }

    Status = Private->BlockIo->WriteBlocks (
                                Private->BlockIo,
                                MediaId,
                                Lba,
                                WorkingBufferSize,
                                WorkingBuffer
                                );

    WorkingBuffer += WorkingBufferSize;
  }

  if (!TransactionComplete && OverRun != 0) {
    //
    // Last bit is not a complete block, so do a read modify write.
    //
    Status = Private->BlockIo->ReadBlocks (
                                Private->BlockIo,
                                MediaId,
                                OverRunLba,
                                Private->BlockSize,
                                Data
                                );

    if (EFI_ERROR (Status)) {
      goto Done;
    }

    EfiCopyMem (Data, WorkingBuffer, OverRun);

    Status = Private->BlockIo->WriteBlocks (
                                Private->BlockIo,
                                MediaId,
                                OverRunLba,
                                Private->BlockSize,
                                Data
                                );
    if (EFI_ERROR (Status)) {
      goto Done;
    }
  }
```

```
Done:
  if (Data) {
    gBS->FreePool (Data);
  }

  return Status;
}
```

D.7 Disk I/O Driver - ComponentName.c

```
/*++

Copyright (c)  1999 - 2003 Intel Corporation. All rights reserved
This software and associated documentation (if any) is furnished
under a license and may only be used or copied in accordance
with the terms of the license. Except as permitted by such
license, no part of this software or documentation may be
reproduced, stored in a retrieval system, or transmitted in any
form or by any means without the express written consent of
Intel Corporation.


Module Name:

  ComponentName.c

Abstract:

--*/

#include "DiskIo.h"

//
// EFI Component Name Functions
//
EFI_STATUS
DiskIoComponentNameGetDriverName (
  IN  EFI_COMPONENT_NAME_PROTOCOL  *This,
  IN  CHAR8                        *Language,
  OUT CHAR16                       **DriverName
  );

EFI_STATUS
DiskIoComponentNameGetControllerName (
  IN  EFI_COMPONENT_NAME_PROTOCOL  *This,
  IN  EFI_HANDLE                   ControllerHandle,
  IN  EFI_HANDLE                   ChildHandle        OPTIONAL,
  IN  CHAR8                        *Language,
  OUT CHAR16                       **ControllerName
  );

//
// EFI Component Name Protocol
//
EFI_COMPONENT_NAME_PROTOCOL gDiskIoComponentName = {
  DiskIoComponentNameGetDriverName,
  DiskIoComponentNameGetControllerName,
```

```
  "eng"
};

static EFI_UNICODE_STRING_TABLE mDiskIoDriverNameTable[] = {
  { "eng", L"Generic Disk I/O Driver" },
  { NULL, NULL }
};

EFI_STATUS
DiskIoComponentNameGetDriverName (
  IN  EFI_COMPONENT_NAME_PROTOCOL  *This,
  IN  CHAR8                        *Language,
  OUT CHAR16                       **DriverName
  )
/*++

  Routine Description:
    Retrieves a Unicode string that is the user readable name of the EFI Driver.

  Arguments:
    This       - A pointer to the EFI_COMPONENT_NAME_PROTOCOL instance.
    Language   - A pointer to a three character ISO 639-2 language identifier.
                 This is the language of the driver name that that the caller
                 is requesting, and it must match one of the languages specified
                 in SupportedLanguages. The number of languages supported by a
                 driver is up to the driver writer.
    DriverName - A pointer to the Unicode string to return. This Unicode string
                 is the name of the driver specified by This in the language
                 specified by Language.

  Returns:
    EFI_SUCCES            - The Unicode string for the Driver specified by This
                            and the language specified by Language was returned
                            in DriverName.
    EFI_INVALID_PARAMETER - Language is NULL.
    EFI_INVALID_PARAMETER - DriverName is NULL.
    EFI_UNSUPPORTED       - The driver specified by This does not support the
                            language specified by Language.

--*/
{
  return EfiLibLookupUnicodeString (
           Language,
           gDiskIoComponentName.SupportedLanguages,
           mDiskIoDriverNameTable,
           DriverName
           );
}

EFI_STATUS
DiskIoComponentNameGetControllerName (
  IN  EFI_COMPONENT_NAME_PROTOCOL  *This,
  IN  EFI_HANDLE                   ControllerHandle,
  IN  EFI_HANDLE                   ChildHandle        OPTIONAL,
  IN  CHAR8                        *Language,
  OUT CHAR16                       **ControllerName
  )
/*++
```

```
  Routine Description:
    Retrieves a Unicode string that is the user readable name of the controller
    that is being managed by an EFI Driver.

  Arguments:
    This             - A pointer to the EFI_COMPONENT_NAME_PROTOCOL instance.
    ControllerHandle - The handle of a controller that the driver specified by
                       This is managing. This handle specifies the controller
                       whose name is to be returned.
    ChildHandle      - The handle of the child controller to retrieve the name
                       of. This is an optional parameter that may be NULL. It
                       will be NULL for device drivers. It will also be NULL
                       for a bus drivers that wish to retrieve the name of the
                       bus controller. It will not be NULL for a bus driver
                       that wishes to retrieve the name of a child controller.
    Language         - A pointer to a three character ISO 639-2 language
                       identifier. This is the language of the controller name
                       that that the caller is requesting, and it must match one
                       of the languages specified in SupportedLanguages. The
                       number of languages supported by a driver is up to the
                       driver writer.
    ControllerName   - A pointer to the Unicode string to return. This Unicode
                       string is the name of the controller specified by
                       ControllerHandle and ChildHandle in the language
specified
                       by Language from the point of view of the driver
specified
                       by This.

  Returns:
    EFI_SUCCESS           - The Unicode string for the user readable name in the
                            language specified by Language for the driver
                            specified by This was returned in DriverName.
    EFI_INVALID_PARAMETER - ControllerHandle is not a valid EFI_HANDLE.
    EFI_INVALID_PARAMETER - ChildHandle is not NULL and it is not a valid
                            EFI_HANDLE.
    EFI_INVALID_PARAMETER - Language is NULL.
    EFI_INVALID_PARAMETER - ControllerName is NULL.
    EFI_UNSUPPORTED       - The driver specified by This is not currently
                            managing the controller specified by
                            ControllerHandle and ChildHandle.
    EFI_UNSUPPORTED       - The driver specified by This does not support the
                            language specified by Language.

--*/
{
  return EFI_UNSUPPORTED;
}
```

Appendix E
EFI 1.10.14.62 Sample Drivers

This appendix lists all the sample drivers that are available in the *EFI Sample Implementation*. Table E-1 shows in which build tips all the drivers are used, Table E-2 describes the property codes used in Table E-3, and Table E-3 shows the properties of each driver.

Table E-1. EFI Driver Build Tips

Driver	BIOS32	IA-32Emb	SAL64	Nt32	Ia32Drivers	IpfDrivers
AtapiPassThru		X			X	X
Bis	X	X		X	X	X
CirrusLogic5430	X	X			X	X
Console\ConPlatform	X	X	X	X	X	X
Console\ConSplitter	X	X	X	X	X	X
Console\GraphicsConsole	X	X	X	X	X	X
Console\Terminal	X	X	X	X	X	X
DebugPort					X	X
DebugSupport					X	X
Decompress	X	X	X	X	X	X
DiskIo	X	X	X	X	X	X
Ebc	X	X	X	X	X	X
FileSystem\Fat	X	X	X	X	X	X
Ide		X			X	X
IsaBus	X	X	X		X	X
IsaFloppy		X			X	X
IsaSerial	X	X	X		X	X
Partition	X	X	X	X	X	X
PcatIsaAcpi		X			X	X
PcatIsaAcpiBios	X		X		X	X
PcatPciRootBridge	X	X	X		X	X
PciBus	X	X	X	X	X	X
PciVgaMiniPort		X			X	X
Ps2Keyboard		X			X	X
Ps2Mouse		X			X	X
PxeBc	X	X	X		X	X
SerialMouse	X	X	X	X	X	X

continued

Table E-1. EFI Driver Build Tips (continued)

Driver	BIOS32	IA-32Emb	SAL64	Nt32	Ia32Drivers	IpfDrivers
Snp32_64	X	X	X		X	X
Undi	X	X	X		X	X
Usb\Uhci		X			X	X
Usb\UsbBot		X			X	X
Usb\UsbBus		X			X	X
Usb\UsbCbi		X			X	X
Usb\UsbKb		X			X	X
Usb\UsbMassStorage		X			X	X
Usb\UsbMouse		X			X	X
VgaClass	X	X	X		X	X
WinNtThunk\BlockIo				X		
WinNtThunk\Console				X		
WinNtThunk\SerialIo				X		
WinNtThunk\SimpleFileSystem				X		
WinNtThunk\Uga				X		
WinNtThunk\WinNtBusDriver				X		
WinNtThunk\WinNtPciRootBridge				X		
BiosInt\BiosKeyboard	X		X			
BiosInt\BiosSnp16	X		X			
BiosInt\BiosVga			X			
BiosInt\BiosVgaMiniPort	X		X			
BiosInt\Disk	X		X			

Table E-2. EFI Driver Property Codes

Field	Field Value	Description
DB		Number of Driver Binding Protocols installed in the driver entry point.
CFG		Y if the Driver Configuration Protocol is installed in the driver entry point.
DIAG		Y if the Driver Diagnostics Protocol is installed in the driver entry point.
CN		Y if the Component Name Protocol is installed in the driver entry point.
Driver Class	Bus	Bus driver.
	Device	Device driver.
	Hybrid	Hybrid driver.
	Root Bridge	Root bridge driver.
	Service	Service driver.
	Init	Initializing driver.

continued

Table E-2. EFI Driver Property Codes (continued)

Field	Field Value	Description
Child	All	All child handles in first call to Start().
	1/ALL	Can create 1 child handle at a time or all child handles in Start().
	1	Creates at most 1 child handle in Start().
	0	Create no child handles in Start(). Used for hot-plug bus types.
Driver Type	BS	EFI Boot Services driver.
	RT	EFI Runtime Services driver.
Unload		Y if the driver is unloadable.

Table E-3. EFI Driver Properties

Driver	DB	CFG	DIAG	CN	Driver Class	Child	Parent	Driver Type	Unload	Hot Plug
AtapiPassThru	1	-	-	Y	Device	-	1	BS	-	-
Bis	-	-	-	-	Service	-	1	BS	-	-
CirrusLogic5430	1	-	-	Y	Device	-	1	BS	-	-
Console\ConPlatform	2	-	-	Y	Device	-	1	BS	-	-
Console\ConSplitter	3	-	-	Y	Bus	ALL	>1	BS	-	-
Console\GraphicsConsole	1	-	-	Y	Device	-	1	BS	-	-
Console\Terminal	1	-	-	Y	Hybrid	1	1	BS	-	-
DebugPort	1	-	-	Y	Bus	1	1	BS	-	-
DebugSupport	-	-	-	-	Service	-	1	BS	-	-
Decompress	-	-	-	-	Service	-	1	BS	-	-
DiskIo	1	-	-	Y	Device	-	1	BS	-	-
Ebc	-	-	-	-	Service	-	1	BS	-	-
FileSystem\Fat	1	-	-	Y	Device	-	1	BS	Y	-
Ide	1	Y	Y	Y	Hybrid	1/ALL	1	BS	-	-
IsaBus	1	-	-	Y	Bus	ALL	1	BS	-	-
IsaFloppy	1	-	-	Y	Device	-	1	BS	-	-
IsaSerial	1	-	-	Y	Bus	1	1	BS	-	-
Partition	1	-	-	Y	Bus	ALL	1	BS	-	-
PcatIsaAcpi	1	-	-	Y	Device	-	1	BS	-	-
PcatIsaAcpiBios	1	-	-	Y	Device	-	1	BS	-	-
PcatPciRootBridge	-	-	-	-	Root Bridge	-	1	BS	-	-

continued

Table E-3. EFI Driver Properties (continued)

Driver	DB	CFG	DIAG	CN	Driver Class	Child	Parent	Driver Type	Unload	Hot Plug
PciBus	1	-	-	Y	Bus	1/ALL	1	BS	-	-
PciVgaMiniPort	1	-	-	Y	Device	-	1	BS	-	-
Ps2Keyboard	1	-	-	Y	Device	-	1	BS	-	-
Ps2Mouse	1	-	-	Y	Device	-	1	BS	-	-
PxeBc	1	-	-	Y	Device	-	1	BS	-	-
PxeDhcp4	1	-	-	Y	Device	-	1	BS	-	-
ScsiBus	1	-	-	Y	Bus	1/ALL	1	BS	-	-
ScsiDisk	1	-	-	Y	Device	-	1	BS	-	-
SerialMouse	1	-	-	Y	Bus	1	1	BS	-	-
Snp32_64	1	-	-	Y	Device	-	1	BS	-	-
Undi	1	-	-	Y	Bus	1	1	RT	-	-
Usb\Uhci	1	-	-	Y	Device	-	1	BS	-	-
Usb\UsbBot	1	-	-	Y	Device	-	1	BS	-	-
Usb\UsbBus	1	-	-	Y	Hybrid	0	1	BS	-	Y
Usb\UsbCbi	1	-	-	Y	Device	-	1	BS	-	-
Usb\UsbKb	1	-	-	Y	Device	-	1	BS	-	-
Usb\UsbMassStorage	1	-	-	Y	Device	-	1	BS	-	-
Usb\UsbMouse	1	-	-	Y	Device	-	1	BS	-	-
VgaClass	1	-	-	Y	Device	-	1	BS	-	-
WinNtThunk\BlockIo	1	-	-	Y	Device	-	1	BS	-	-
WinNtThunk\Console	1	-	-	Y	Device	-	1	BS	-	-
WinNtThunk\SerialIo	1	-	-	Y	Bus	1	1	BS	-	-
WinNtThunk\SimpleFileSystem	1	-	-	Y	Device	-	1	BS	-	-
WinNtThunk\Uga	1	-	-	Y	Device	-	1	BS	-	-
WinNtThunk\WinNtBusDriver	1	-	-	Y	Hybrid	1/ALL	1	BS	-	-
WinNtThunk\WinNtPciRootBridge	-	-	-	-	Root Bridge	-	1	BS	-	-
BiosInt\BiosKeyboard	1	-	-	Y	Device	-	1	BS	-	-
BiosInt\BiosSnp16	1	-	-	Y	Bus	1	1	BS	-	-
BiosInt\BiosVga	1	-	-	Y	Device	-	1	BS	-	-
BiosInt\BiosVgaMiniPort	1	-	-	Y	Device	-	1	BS	-	-
BiosInt\Disk	-	-	-	-	Root Bridge	-	1	BS	-	-

Appendix F
Glossary

Table F-1 defines terms that are used in this document. See the glossary in the *EFI 1.10 Specification* for definitions of additional terms.

Table F-1. Definitions of Terms

Term	Definition
<Enumerated Type>	Element of an enumeration. Type INTN.
ACPI	Advanced Configuration and Power Interface.
ANSI	American National Standards Institute.
API	Application programming interface.
ASCII	American Standard Code for Information Interchange.
ATAPI	Advanced Technology Attachment Packet Interface.
BAR	Base Address Register.
BBS	*BIOS Boot Specification.*
BC	Base Code.
BEV	Bootstrap Entry Vector. A pointer that points to code inside an option ROM that will directly load an OS.
BIOS	Basic input/output system.
BIS	Boot Integrity Services.
BM	Boot manager.
BOOLEAN	Logical Boolean. 1-byte value containing a 0 for **FALSE** or a 1 for **TRUE**. Other values are undefined.
BOT	Bulk-Only Transport.
BS	EFI Boot Services Table or EFI Boot Service(s).
CBI	Control/Bulk/Interrupt Transport.
CBW	Command Block Wrapper.
CHAR16	2-byte character. Unless otherwise specified, all strings are stored in the UTF-16 encoding format as defined by Unicode 2.1 and ISO/IEC 10646 standards.
CHAR8	1-byte character.
CID	Compatible ID.
CONST	Declares a variable to be of type const. This modifier is a hint to the compiler to enable optimization and stronger type checking at compile time.
CR	Containing Record.
CRC	Cyclic Redundancy Check.

continued

Table F-1. Definitions of Terms (continued)

Term	Definition
CSW	Command Status Wrapper.
DAC	Dual Address Cycle.
DHCP4	Dynamic Host Configuration Protocol Version 4.
DID	Device ID.
DIG64	*Developer's Interface Guide for 64-bit Intel Architecture-based Servers.*
DMA	Direct Memory Access.
EBC	EFI Byte Code.
ECR	Engineering Change Request.
EFI	Extensible Firmware Interface.
EFI_EVENT	Handle to an event structure. Type VOID *.
EFI_GUID	128-bit buffer containing a unique identifier value. Unless otherwise specified, aligned on a 64-bit boundary.
EFI_HANDLE	A collection of related interfaces. Type VOID *.
EFI_IP_ADDRESS	16-byte buffer aligned on a 4-byte boundary. An IPv4 or IPv6 internet protocol address.
EFI_Ipv4_ADDRESS	4-byte buffer. An IPv4 internet protocol address.
EFI_Ipv6_ADDRESS	16-byte buffer. An IPv6 internet protocol address.
EFI_LBA	Logical block address. Type UINT64.
EFI_MAC_ADDRESS	32-byte buffer containing a network Media Access Controller address.
EFI_STATUS	Status code. Type INTN.
EFI_TPL	Task priority level. Type UINTN.
EISA	Extended Industry Standard Architecture.
FAT	File allocation table.
FIFO	First In First Out.
FPSWA	Floating Point Software Assist.
FRU	Field Replaceable Unit.
FTP	File Transfer Protocol.
GPT	Guided Partition Table.
GUID	Globally Unique Identifier.
HC	Host controller.
HID	Hardware ID.
I/O	Input/output.
IA-32	32-bit Intel architecture.
IBV	Independent BIOS vendor.
IDE	Integrated Drive Electronics.

continued

Table F-1. Definitions of Terms (continued)

Term	Definition
IEC	International Electrotechnical Commission.
IHV	Independent hardware vendor.
IN	Datum is passed to the function.
INT	Interrupt.
INT16	2-byte signed value.
INT32	4-byte signed value.
INT64	8-byte signed value.
INT8	1-byte signed value.
INTN	Signed value of native width. (4 bytes on IA-32, 8 bytes on Itanium architecture operations)
IPF	Itanium processor family.
Ipv4	Internet Protocol Version 4.
Ipv6	Internet Protocol Version 6.
ISA	Industry Standard Architecture.
ISO	Industry Standards Organization.
iSCSI	SCSI protocol over TCP/IP.
KB	Keyboard.
LAN	Local area network.
LUN	Logical Unit Number.
MAC	Media Access Controller.
MMIO	Memory Mapped I/O.
NIC	Network interface controller.
NII	Network Interface Identifier.
NVRAM	Nonvolatile RAM.
OEM	Original equipment manufacturer.
OHCI	Open Host Controller Interface.
OpROM	Option ROM.
OPTIONAL	Datum that is passed to the function is optional, and a **NULL** may be passed if the value is not supplied.
OS	Operating system.
OUT	Datum is returned from the function.
PCI	Peripheral Component Interconnect.
PCMCIA	Personal Computer Memory Card International Association.
PE	Portable Executable.
PE/COFF	PE32, PE32+, or Common Object File Format.
PNPID	Plug and Play ID.

continued

Table F-1. Definitions of Terms (continued)

Term	Definition
POST	Power On Self Test.
PPP	Point-to-Point Protocol.
PUN	Physical Unit Number.
PXE	Preboot Execution Environment.
PXE BC (or PxeBc)	PXE Base Code Protocol.
QH	Queue Head.
RAID	Redundant Array of Inexpensive Disks.
RAM	Random access memory.
ROM	Read-only memory.
RT	EFI Runtime Table and EFI Runtime Service(s).
SAL	System Abstraction Layer.
SCSI	Small Computer System Interface.
SIG	Special Interest Group.
S.M.A.R.T.	Self-Monitoring Analysis Reporting Technology.
SMBIOS	System Management BIOS.
SMBus	System Management Bus.
SNP	Simple Network Protocol.
SPT	SCSI Pass Thru.
ST	EFI System Table
STATIC	The function has local scope. This modifier replaces the standard C static key word, so it can be overloaded for debugging.
TCP/IP	Transmission Control Protocol/Internet Protocol.
TD	Transfer Descriptor.
TPL	Task Priority Level.
UART	Universal Asynchronous Receiver-Transmitter.
UGA	Universal Graphics Adapter.
UHCI	Universal Host Controller Interface.
UID	Unique ID.
UINT16	2-byte unsigned value.
UINT32	4-byte unsigned value.
UINT64	8-byte unsigned value.
UINT8	1-byte unsigned value.
UINTN	Unsigned value of native width. (4 bytes on IA-32, 8 bytes on Itanium architecture operations)
UNDI	Universal Network Driver Interface.

continued

Table F-1. Definitions of Terms (continued)

Term	Definition
USB	Universal Serial Bus.
USBSTS	Status register in the USB host controller.
VGA	Video graphics array.
VID	Vendor ID.
VOID	Undeclared type.
VOLATILE	Declares a variable to be volatile and thus exempt from optimization to remove redundant or unneeded accesses. Any variable that represents a hardware device should be declared as `VOLATILE`.
VT100	A terminal and serial protocol originally defined by Digital Equipment Corporation. Limited to 7-bit ASCII.
VT-UTF8	A serial protocol definition that extends VT-100 to support Unicode. See ISO standard 10646-1: 2000 for more information.