

with HP ProLiant Gen9 Server Blades

Table of contents

Introduction	2
Why UEFI now?	2
UEFI Systems	4
Role of Virtual Connect	4
How does Virtual Connect set the server profile?	4
Role of VC in UEFI supported servers	6
How does VC assign boot modes?	7
PXE IP boot ordering	8
Conclusion	9
Support and other resources	10
Resources	10
Before you contact HP	10
HP contact information	10
Documentation feedback	11
Version	11
Acronyms and abbreviations	12

This document briefly describes UEFI and how Virtual Connect supports UEFI through its server profiles.

Introduction

The use of the BIOS for PCs and servers extends over three decades. The BIOS runs on many PC and server systems today and is based on code initially developed in 1979. The UEFI addresses the need for a new firmware interface that is designed for the operating system and hardware architectures of today, where PCs and servers commonly have one or more multi-core processors, numerous I/O ports of different types, PCI devices with their own controllers, and terabytes of memory and storage space. This level of hardware complexity requires constant ad hoc fixes to legacy BIOS. Patching a system not originally designed for these expansive systems makes such fixes burdensome for operations and management.

UEFI is a firmware interface with the same basic purpose of legacy BIOS, but it uses a different firmware stack. UEFI employs a modular structure of drivers and files grouped into services that adapts well to complex and varied architectures.

UEFI's framework of coding is designed to scale with server architectural trends, including new processors. UEFI encourages architectural freedom, and it is agnostic to the type of processor used in a platform, such as energy-efficient Reduced Instruction Set Computing (RISC) processors.

	Legacy BIOS	UEFI Firmware
Architecture	x86/x64 only	Agnostic (x86, x64, ARM, ia64)
Programming language	Assembly	с
Processors supported	Intel Architecture	Intel Architecture, Itanium, and ARM officially supported
Processor mode used	16-bit, single CPU, single-threaded	Native 64- or 32-bit, single CPU, single threaded
Expansion card firmware	Legacy Option ROMs	UEFI drivers or option ROMs
Provided services	Interrupts	Protocols
Video support	Int10hVBIOSVGA	Graphics Output Protocol (GOP)
Storage support	 Int13h Master Boot Record (MBR) Partitioning 	 Block I/O protocol GUID Partition Table (GPT) Master Boot Record (MBR) partitioning
Peripheral and feature setup	Function key setupNo industry standard	UEFI Human Interface Infrastructure (HII) Protocol as industry standard.
OS boot loader	 Int19h loads 16-bit boot sector in MBR One boot loader per device 	 UEFI loads boot laoder executable files per priority/ordering defined by UEFI specficiation. Multiple boot loader files, unique names or paths can coexist on the same partition or device.
Mode	16-bit (real mode)	32/64 bit
OS handoff	No clear definition	ExitBootServices() function defined by UEFI specification.

Advantages of UEFI over BIOS:

UEFI was introduced to HP Proliant Gen9 Server Blades (x86 based), which means Virtual Connect now supports UEFI and all of its features.

Why UEFI now?

HP has been involved in the UEFI forum since its inception in 2005, and is currently on the board of directors with over 250 members. We implemented EFI, and then UEFI in our HP Integrity server line, and have been using UEFI in our client desktops, notebooks, and printers. Recent trends have prompted a broad UEFI application to the enterprise server industry, including:

- Expanded boot volume capacity
- Pre-boot security requirements
- Pre-operating system networking needs
- Pre-boot manageability needs

- UEFI-only operating system requirements
- Expanded amount of I/O ports and bootable devices
- Remote provisioning



Based on the rapid deployment of UEFI as an industry standard BIOS, HP has released different products that support UEFI. Below are the list of products that HP has developed, or is developing, with UEFI as its core BIOS stack.

UEFI - Converged Firmware Infrastructure



Since the transition to UEFI must be gradual, HP has created a range of servers to smoothly transition customers from legacy BIOS to UEFI based servers. Different types of servers are available based on supported boot modes.

UEFI Systems

UEFI Systems Classes



- Class 2 System: UEFI definition of a system that can boot into UEFI mode or Legacy BIOS mode
- Class 3 System: UEFI definition of a system that can only boot into Native UEFI mode
- CSM: Compatibility Support Module. Allows as Class 2 UEFI system to boot into Legacy BIOS mode (Emulation)

Now that the HP server blades are available with UEFI support, it's the right time for the other modules which coexist with the servers in the C7000 or C3000 enclosure to support UEFI, so that the customers are not faced with any issues in utilizing these extended features of UEFI. Leading the support of UEFI is Virtual Connect, which plays an important role of virtualizing the server blade's SAN and LAN connections to the outside network.

Role of Virtual Connect

HP Virtual Connect for HP Blade System simplifies the setup of server connections to LANs and SANs, allowing server administrators to quickly add or replace servers and move workloads without needing to involve network and storage teams. HP Virtual Connect Flex-10 is the first interconnect technology to reduce the cost and amount of network equipment needed compared to switches, while providing precise bandwidth control for every server Ethernet connection. Virtual Connect Enterprise Manager provides a single intuitive console that manages server connectivity for hundreds of Virtual Connect domains and thousands of servers.

All you need to do is have your system, LAN, and SAN administrators provision your LAN and SAN connections in Virtual Connect. The next time you add a server, your system administrator won't have to call the LAN or SAN administrators — the system administrator can add the new server blade into the enclosure, assign the appropriate server profile containing your defined network connections, and go! There is no need to coordinate with the LAN and SAN administrators.

How does Virtual Connect set the server profile?

Configuring a server profile in Virtual Connect depends on 2 modules:

- 1. The Enclosure Manager, commonly called Onboard Administrator or OA
- 2. The server blade iLO (Integrated Lights-Out)

The server profile is typically configured in the following order:

- 1. The server profile is configured using the VC Manager (VCM) GUI or CLI.
- 2. The server blade is installed.
- 3. VCM detects the server blade is installed, and then reads the FRU (Field Replaceable Unit) data for each interface, like the HBA, the NIC, etc.
- 4. VCM writes the server profile information to the server.
- 5. When the server is powered up, the CPU, BIOS, and NIC/HBA option ROM software writes the server profile information to the EFI (Extensible Firmware interface).

- 6. A successful write is communicated to VCM via OA.
- 7. The server boots using the configured server profile.



Role of VC in UEFI supported servers

Virtual Connect (VC) allows server administrators to configure the server boot mode in a server profile. HP ProLiant Gen9 server blades use the Unified Extensible Firmware Interface (UEFI) between the server blade firmware and the Operating System (OS). This requires VCM to make changes in how VC configures the boot options for a server blade.

The Virtual Connect UEFI Boot Order support feature has made major changes to our processing of physical servers especially regarding the boot mode configuration. There are changes to determine if a server is UEFI capable and a Class 2 or Class 3 UEFI system. The handling of a Class 2 UEFI capable system the most complex.



There is a new option on the server profile to specify if a server should boot in UEFI mode or in Legacy mode. This setting is not used when the profile is assigned to a non-UEFI capable server. If the profile is assigned to a Class 2 system, VC will have to change the boot mode to match the profile setting. A configuration error occurs if a profile in Legacy boot mode is assigned to a Class 3 UEFI system.

When a server is installed, and then added to VCM, VC checks to see if the server is UEFI capable or not. If the physical server is a Class 2 system, then VC needs to read and potentially change the boot mode to match that of the profile as shown in the flowchart above.

The table below shows the profile boot mode and the current boot mode of the server. The action to take indicates what happens based on the Customer's choice. The requested boot mode reflects the boot that needs to be set.

Profile Boot Mode	Current Boot Mode	Requested Boot Mode	Action to take
Legacy	Legacy	_	No action
Legacy	UEFI	Legacy	Set requested boot mode, and then reboot
UEFI	Legacy	UEFI	Set requested boot mode, and then reboot
UEFI	UEFI	—	No action

If the physical server is not UEFI capable VC will ignore the profile boot mode. An informational status will indicate that this profile has a configuration that doesn't match the server it is assigned to.

How does VC assign boot modes?

As part of processing a newly installed server, VC reads information from the system BIOS to determine if the server is UEFI capable or not. This information is stored until the profile is assigned. A new server profile setting is added to let the customer decide if a server should boot in UEFI mode or Legacy mode. This setting is only applicable for UEFI servers.

The basis on which the Customer decides whether to boot in UEFI or in Legacy Boot Mode can be any of the below following reasons:

- Functionality requiring UEFI
- Up gradation of existing resources
- OS installation
- Boot Considerations
- Customer needs in the data center

Once the profile is applied after choosing the boot mode, VC writes the boot order EVs (Environment Variables) while the server is paused from booting. This happens after the POST (Power-On Self-Test) is complete.

On Legacy BIOS based systems, remote management software polls the server for an End of POST (EOP) event as an indicator that the BIOS has finished going through POST. This is necessary to safely read current BIOS configuration EVs, and then write them. A typical scenario would be to change the BIOS boot order after an EOP, and then reboot the server.

HP ProLiant Gen8 servers support a special EV used by Virtual Connect to instruct the BIOS to delay for a specified number of seconds after EOP has been signaled, but before rebooting the server. This allows software enough time to read and modify the BIOS settings, including changing the boot order.

The following diagram illustrates the boot sequence on HP ProLiant Gen8 servers.



RBSU is a configuration utility embedded in HP ProLiant servers that performs a wide range of configuration activities that include the following:

- Configuring system devices and installed options
- Enabling and disabling system features
- Displaying system information
- Selecting the primary boot controller
- Configuring memory options
- Language selection.

On UEFI based servers, the End of POST does not occur until a real OS is booted. This does not happen on bare metal systems, where a production or service OS has not been installed yet.

Software that needs to poll End of POST will have to instead use a new signal called POST Discovery Complete.

The following diagram illustrates the POST sequence on UEFI based systems:



Software can NOT change EVs

In case of Legacy Boot Mode when the server boots and reaches POST complete, VCM will read the current bootable devices from an EV and re-order them based on the profile configuration. The server then reboots with the newly configured boot ordering.

In case of UEFI Boot Mode when the server reaches discovery complete, the design goal is that VC sets an EV. The EV instructs BIOS to boot through Discovery complete, issue a POST complete, then wait for some time and reboot. This allows VC to write the EVs at that time. VC doesn't actually read any EVs in UEFI mode. VC uses the OA and SMBIOS info and VC profile to determine which devices are bootable. It then configures those in the desired boot order EVs.

To support UEFI Boot Order specification, VC includes Boot mode pull down option to the profile define/edit screen. The boot mode options are as follows:

- Auto (default)
- Legacy (Legacy BIOS)
- UEFI

To support this specification, VC-GUI has a drop down (UEFI boot system) in profile screen. Auto is the default boot mode. Users with server privilege can add/change UEFI boot mode.

Profile					
Profile Name Profile1	Network Access Drost Foor Hode	Mide Unused Flexibles Advanced Profile Settings			
Ethemet Ada	opter Connections UEFI				
Port Network	Lepky Street Lepky	PRE/IP Book Crober	Walkcard Filter	MAC	Adam
1 Unassign	wd O PREFERRED	USE-BIOS	None	VC-DEFINED	
2 Unassign	ed O PREFERRED	USE-BIOS	None	VC-DEFINED	
+ Add					
ISCSI HBA C	onnections				
Port Network	Name Status Port S	peed Type East	Setting	MAC .	Action
+ Add					
FC HBA Con	nections				
Port Connecte	d To PC SAN Name	Sister	Port Speed Type	1000Pk	Adan
1 Bay 5	Unassigned	0	AUTO	VC-DEFINED	
2 Bay 6	Unassigned	0	AUTO	VC-DEFINED	Delete
First Channel Bool Parameters					
FCoE HBA Connections					
Port Connected	To PC SAN / PCsE Network Name Ty	pe Status PortSpeed Type	101076	MAC	Action
1 Bay 1	Unassigned	PREFERRED	VC-DEFINED	VC-DEFINED	

PXE IP boot ordering

PXE works with the NIC by making it function like a boot device. When a PXE-enabled NIC sends out a DHCP broadcast request, the DHCP server responds with boot information for the client.

The following steps describe how this works using a Linux boot image:

- 1. A Target Machine (either bare metal or with boot sector removed) boots.
- 2. The PXE-enabled NIC broadcasts a DHCP request.
- 3. A DHCP server receives the request, and then responds to the client, offering the following information:
 - An IP address, subnet mask, DNS, etc.
 - The address of the TFTP server
 - The location of boot files on the TFTP server
- 4. When the client receives this information, it contacts the TFTP server to get the boot image.
- 5. The TFTP server sends the boot image (pxelinux.0), and then the client executes it.
- 6. The Operating System is loaded.

By default, the boot image searches for the pxelinux.cfg directory on the TFTP server for boot configuration files.

UEFI allows PXE-enabled NICs to boot on IPv6 systems. UEFI capable systems also have the option to specify the PXE IP boot ordering. The VC GUI updates the Ethernet connection configuration of a profile with this setting. The user can decide the order in which IPv4 and/or IPv6 are used for PXE booting. This setting is configured to the server during the profile assignment. To support this specification, VC provides a PXE/IP Boot Order drop-down option to specify the PXE/IP boot ordering:

- Auto (It will be a default PXE IP boot order)
- IPv40nly
- IPv60nly
- IPv4ThenIPv6
- IPv6ThenIPv4.

Profile							
Profile Name Profile1		Network Access Group Ever Hede	Hide Enused FlexNEs Advanced Profile Settings				
Ethernet A	dapter	Connections					
Fort Netwo	ot liane	Status Port Speed Type	PXE/P Boat Crd	-	Walticest Filter	MAC .	Active
1 Unass	signed	PREFERRED	USE-BIOS		None	VC-DEFINED	
2 Unass	signed	O PREFERRED	USE-EROS	-	None	VC-DEFINED	
+ Add			USE-BIOS DISABLED ENABLED	Auto			
ISCSI HBA	Conne	ections		IPv4	iOnly iOnly		
Fot Netwo	os Name	Status Port Speed	Type Book Setting	1Pv1	ThesiPv6	MAC .	Action
+ Add	onnect	1075					
Test Course		PT 510 Name	_		Part Scent Tune	20070	1000
1 Bay 5		Unassigned		0	AUTO	VC-DEFINED	
2 Bay 6		Unassigned		0	AUTO	VC-DEFINED	Dente
Add Three Channel Bast Reconstrue							
FCOE HBA	Conn	ections					
Part Conner	cled Te	PC SAN / PCsE Network Name Type	Status Port Speed Type		877498	UAC .	Adae
1 Bay 1		Unassigned	PREFERRED		VC-DEFINED	VC-DEFINED	-

Conclusion

Our leadership role in the UEFI forum and close association with UEFI development allows us to develop solutions that offer advanced UEFI functionality. The advanced UEFI functionality we have integrated into the Virtual Connect server profile provides a highly manageable and easily provisioned system that helps customers to utilize the maximum benefits of UEFI through Virtual Connect in a Blade System environment.

Learn more at hp.com/go/vc/manuals

Support and other resources

Resources

- hp.com/go/vc/manuals
- hp.com/go/virtualconnect
- hp.com/go/bladesystem

Before you contact HP

Be sure to have the following information available before you call HP:

• Active Health System log (HP ProLiant Gen8 or later products)

Download and have available an Active Health System log for 7 days before the failure was detected. For more information, see the HP iLO 4 User Guide or HP Intelligent Provisioning User Guide on the HP website (http://www.hp.com/go/ilo/docs).

• Onboard Administrator SHOW ALL report (for HP BladeSystem products only)

For more information on obtaining the Onboard Administrator SHOW ALL report, see the HP website (http://www.hp.com/go/OAlog).

- Technical support registration number (if applicable)
- Product serial number
- Product model name and number
- Product identification number
- Applicable error messages
- Add-on boards or hardware
- Third-party hardware or software
- Operating system type and revision level

HP contact information

For United States and worldwide contact information, see the Contact HP website (http://www.hp.com/go/assistance).

In the United States:

- To contact HP by phone, call 1-800-334-5144. For continuous quality improvement, calls may be recorded or monitored.
- If you have purchased a Care Pack (service upgrade), see the Support & Drivers website (http://www8.hp.com/us/en/support-drivers.html). If the problem cannot be resolved at the website, call 1-800-633-3600. For more information about Care Packs, see the HP website (http://pro-aqsama.houston.hp.com/services/cache/10950-0-0-225-121.html).

On a best effort basis only, HP offers technical assistance on low-latency tuning to customers who have followed this guide and still have questions. For more information, contact HP (mailto:low.latency@hp.com). Provide your name and your company's name, as well as your mailing address and your HP contact's name.

Documentation feedback

HP is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback (<u>mailto:docsfeedback@hp.com</u>). Include the document title and part number, version number, or the URL when submitting your feedback.



798696-001 Edition 1 September 2014

Acronyms and abbreviations

BIOS Basic Input Output System DHCP Dynamic Host Configuration Protocol DNS domain name system ΕV environment variable FRU field replaceable unit HBA host bus adapter iLO Integrated Lights-Out 0A Onboard Administrator OPROM Optional ROM in HBA/NIC PXE preboot execution environment RBSU ROM-Based Setup Utility RISC Reduced Instruction Set Computing TFTP Trivial File Transfer Protocol UEFI Unified Extensible Firmware Interface