

Enabling iSCSI boot under UEFI boot mode on Dell 13th generation of PowerEdge servers

A technical white paper detailing the configuration options and process for iSCSI boot for Dell 13th generation of PowerEdge servers.

Version 2.0
March, 2016

Xi Li, Software Dev Staff Engineer
Paul Rubin, Lifecycle Controller Product Manager

Dell | Enterprise Solutions Group

Contents

Executive summary	3
Introduction	4
Network Boot - PXE and iSCSI	4
iSCSI boot from NIC firmware and iSCSI boot from BIOS	4
Configuring iSCSI boot parameters	7
iSCSI Boot Workflow	11
Current Limitations	15
Conclusion	15
Appendix - Acronyms and Terminology	16

Figures

Figure 1. Configuring option-ROM based legacy mode iSCSI boot	5
Figure 2. Accessing Network Settings option from PowerEdge UEFI BIOS	6
Figure 3. Accessing iSCSI Device1 Settings from Network Settings	7
Figure 4. Accessing Connection 1 Settings option from iSCSI Device1 Settings	8
Figure 5. iSCSI boot configuration page one	9
Figure 6. iSCSI boot configuration page two	10
Figure 7. iSCSI boot work flowchart (1)	13
Figure 8. iSCSI boot work flowchart (2)	14
Figure 9. iSCSI boot work flowchart (3)	14

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

Copyright © 2016 Dell Inc. All rights reserved. Dell and the Dell logo are trademarks of Dell Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

March 2016 | Version 2.0

Executive summary

With the growing use of network-based server deployment processes, compute grids and software-defined services on modern servers require multiple options to enable server network boot. Dell 13th generation of PowerEdge servers has recently introduced UEFI-based Internet Small Computer System Interface (iSCSI) boot capabilities that support a range of server deployment and startup scenarios.

This document describes the following aspects for iSCSI boot support configuration within a 13th generation of PowerEdge server BIOS:

- Server network booting; understanding PXE boot and iSCSI boot
- Network boot under legacy BIOS and UEFI modes
- Configuration of the iSCSI boot feature
- The iSCSI boot work flow
- Current known limitations of the iSCSI boot feature

Note: iSCSI solutions rely on multiple datacenter infrastructure components such as network switches, storage subsystems, and server operating systems. For information about other iSCSI solution components, see the appropriate documentation.

Introduction

Network boot is the process of booting a system with an operating image available in a network service. By using network boot, you can reduce the overall data center storage needs, enhance solution reliability through multiple remote boot services, and reduce the recovery time for storage failures.

In the past, Dell PowerEdge servers provided a network boot feature that used the legacy boot mode of individual network interface controllers. This resulted in complexities and performance challenges.

Dell's 13th generation of PowerEdge servers provide an enhanced network boot support by utilizing the UEFI boot environment to implement both PXE and iSCSI-based network boot.

Network Boot - PXE and iSCSI

The following are the two commonly used network boot methodologies in the industry:

- PXE boot - Uses a program download mechanism through a network interface to download and execute a program without using the local disk storage space.
- iSCSI boot - Uses a mechanism called Disk Emulation. Instead of downloading an executable program through a network interface, iSCSI clients or initiators send SCSI commands to a SCSI storage device or target on a remote server and treat the remote storage device as a locally attached storage. iSCSI interactions are used to load an executable program such as an operating system to the memory of the server.

To provide device-independent support for network interfaces from multiple vendors, 13th generation of PowerEdge server BIOS has recently added an iSCSI boot feature in UEFI boot mode.

iSCSI boot from NIC firmware and iSCSI boot from BIOS

Prior to the availability of the PowerEdge UEFI iSCSI boot feature, the iSCSI boot feature was device and vendor specific. The feature was based on the invocation of an iSCSI boot stack that is provided by the PCI option ROM of the network interface.

This form of iSCSI boot is supported on the legacy BIOS mode of the server and it depends on the various network interfaces to implement the iSCSI boot option ROM code.

Figure 1 illustrates the configuration of a device specific iSCSI boot function.

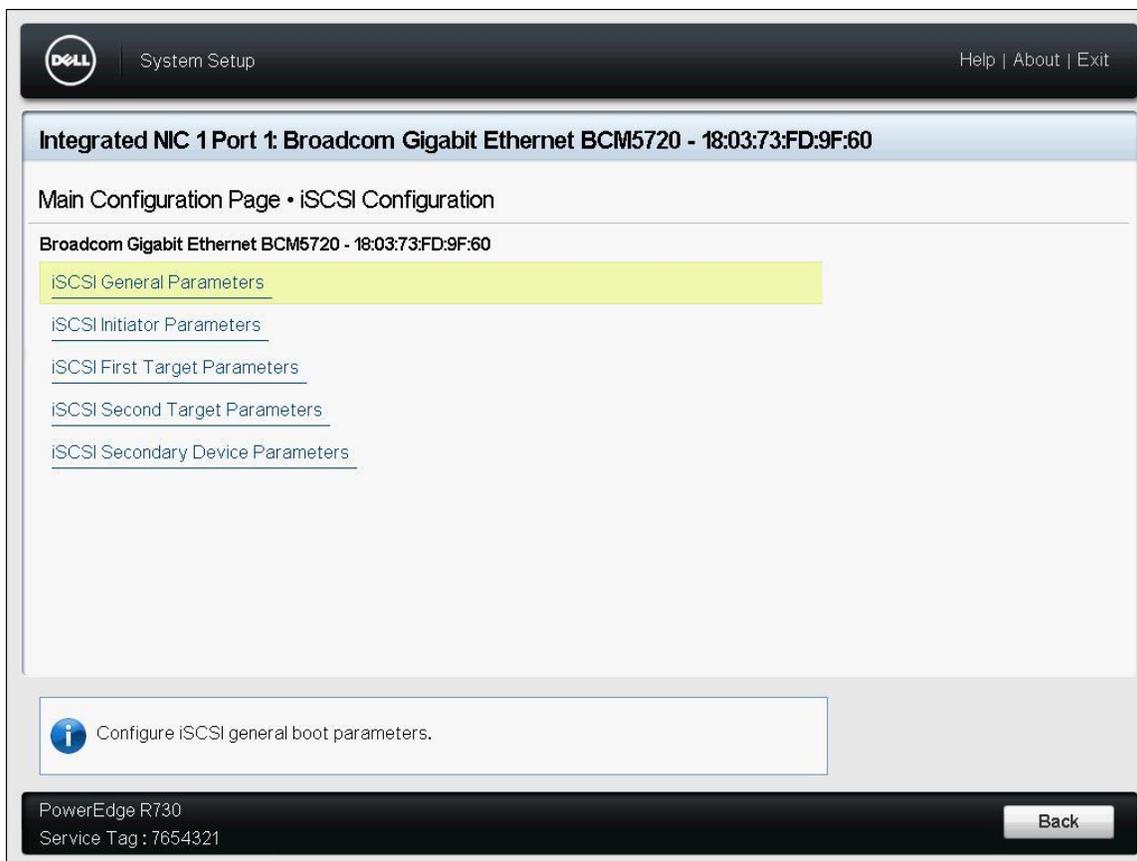


Figure 1. Configuring option-ROM based legacy mode iSCSI boot

The 13th generation of PowerEdge servers UEFI iSCSI boot feature enables efficient multi-vendor support of iSCSI boot by implementing the iSCSI boot function within the PowerEdge UEFI BIOS and by removing the dependency on vendor-specific option ROM support.

The UEFI iSCSI boot feature can be configured through the 13th generation of PowerEdge BIOS user interface or through the PowerEdge iDRAC with Lifecycle Controller command line interface. This white paper illustrates the UEFI iSCSI boot feature using the BIOS user interface.

Assuming a 13th generation of PowerEdge BIOS configured for UEFI boot mode, a **Network Settings** menu (Figure 2) is provided from the **System BIOS** page user interface, enabling the configuration of PXE and iSCSI boot using BIOS-based drivers.

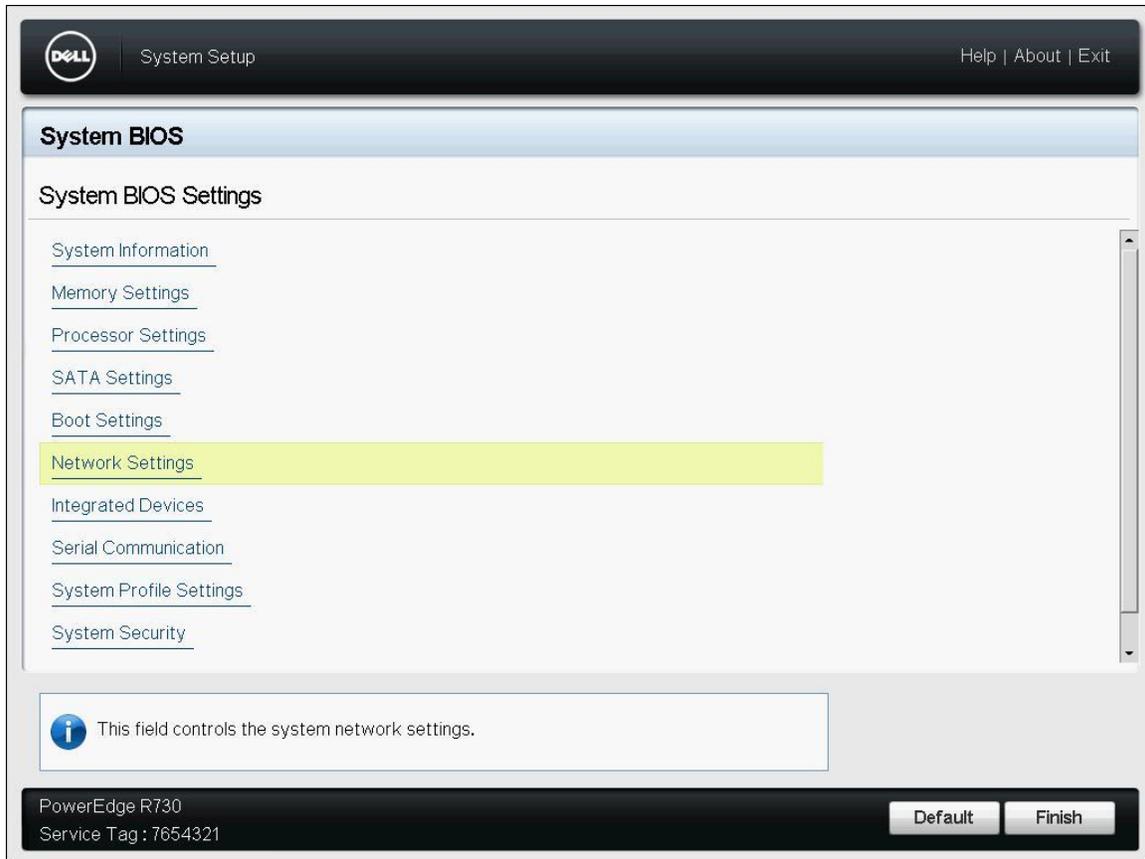


Figure 2. Accessing Network Settings option from PowerEdge UEFI BIOS

Configuring iSCSI boot parameters

To configure the iSCSI boot parameters, perform the following steps:

1. On the **Network Settings** page, enter the initiator name in the **iSCSI Initiator Name** field.

Note: The iSCSI initiator name must follow IQN format, which includes the "iqn." string, date in yyyy-mm format, a period ("."), the reversed domain name of the naming authority, and an optional string within the character set and length boundaries prefixed with a colon (":").

2. Select **Enabled** from the **iSCSI Device1** options.

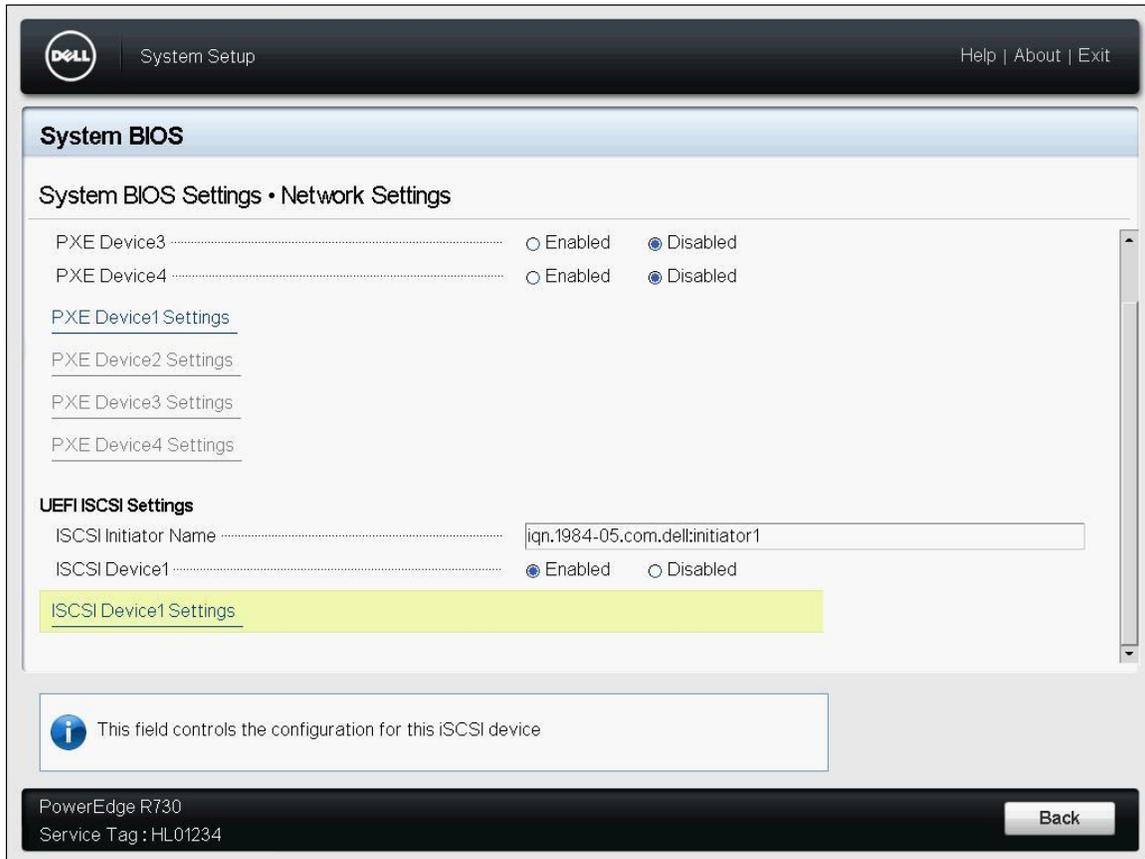


Figure 3. Accessing iSCSI Device1 Settings from Network Settings

3. Click the **iSCSI Device1 Settings** link.
4. Enable one of the two connections – **Connection 1 Settings** or **Connection 2 Settings** as required. For this example, click the **Connection 1 Settings** link.

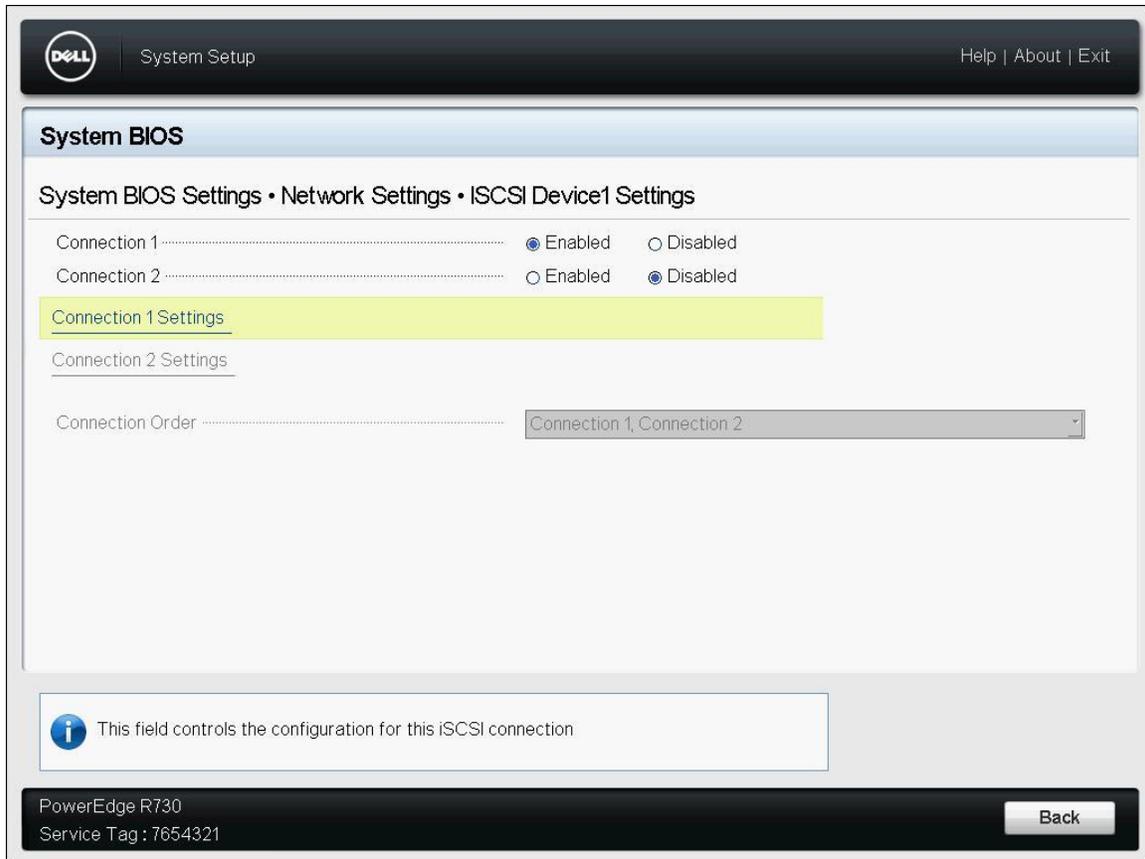


Figure 4. Accessing Connection 1 Settings option from iSCSI Device1 Settings

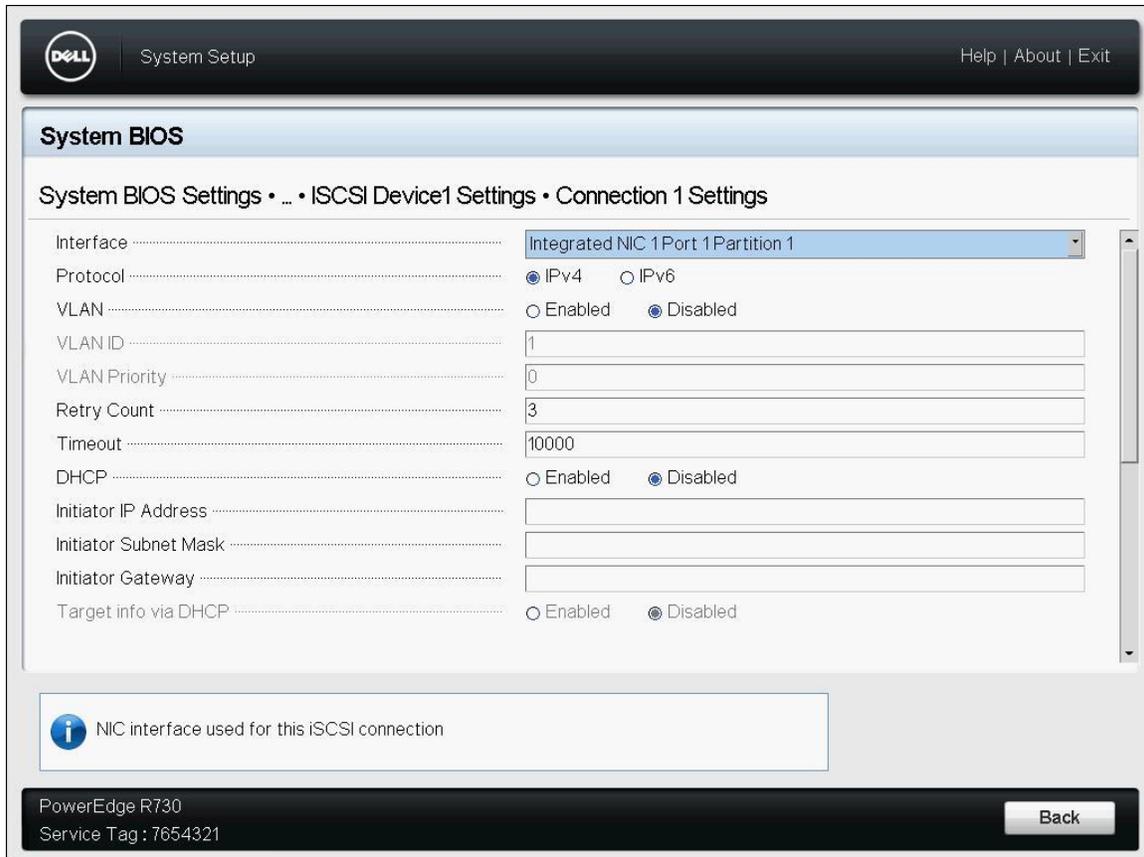


Figure 5. iSCSI boot configuration page one

5. Select the network interface to be configured to perform iSCSI boot from the **Interface** drop-down menu
6. Choose the **IPv4** or **IPv6** IP protocol as appropriate, and if required, use **VLAN**, **VLAN ID** and **VLAN Priority** to configure VLAN.
7. Enter appropriate values in the **Retry Count** and **Timeout** fields respectively to manage interface error and timeout conditions.
8. Select **Enabled** or **Disabled** from the **DHCP** option.

Note: The **DHCP** option is used to select manual or DHCP server-based configuration of the initiator IP, subnet, and gateway addresses. If DHCP is enabled, it is not required to configure the three initiator parameters. If DHCP is disabled, you must manually enter the three initiator parameters.

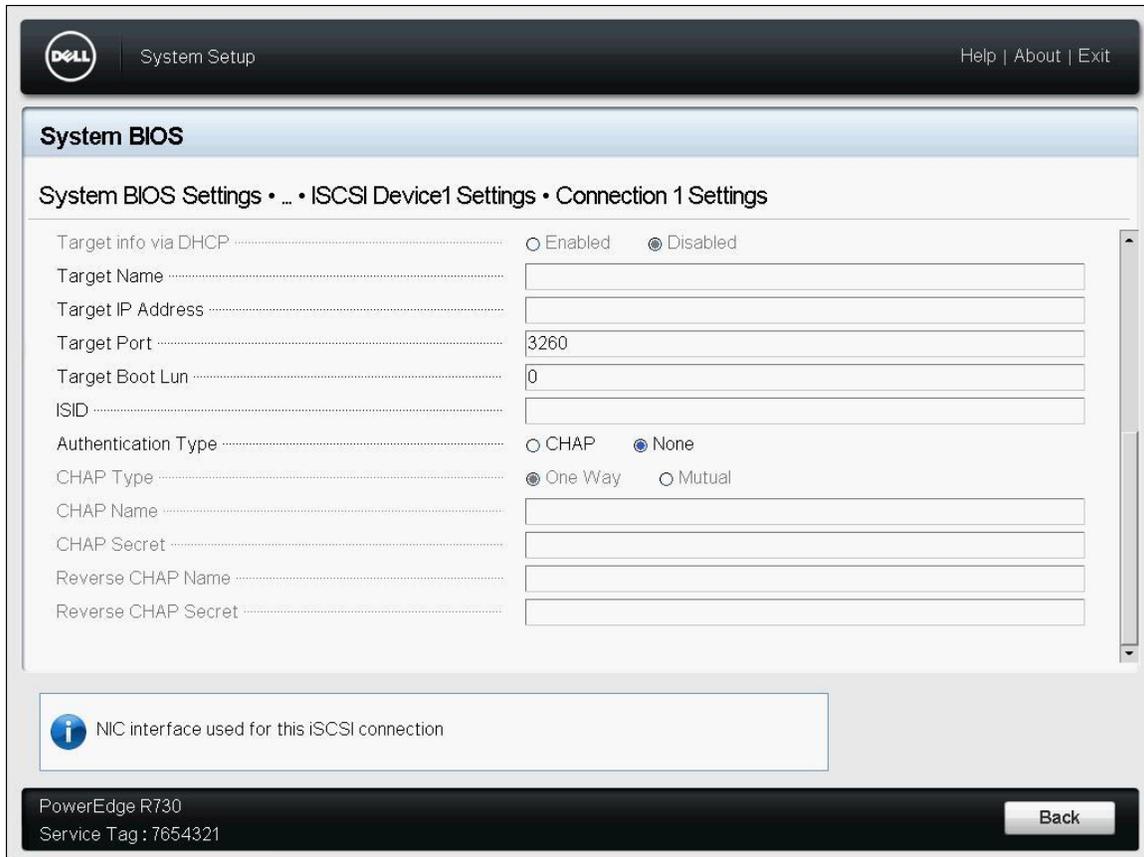


Figure 6. iSCSI boot configuration page two

Similarly, the iSCSI target can be configured through a DHCP server or manually. If the **DHCP** option in the previous step is disabled, the **Target info via DHCP** automatically selects **Disabled**.

9. Enter appropriate values in the **Target Name**, **Target Name IP Address**, **Target Port**, and **Target Boot LUN** fields respectively for static configuration. If **Target info via DHCP** is set to **Enabled**, the information is obtained from the DHCP sever through DHCP option 17, the root path option.

Target Port is the TCP port that iSCSI uses for communication. The **Target Port** is set to 3260 by default. **Target Boot Lun** is the LUN number from which the iSCSI initiator boots.

If required, use the **ISID** option to set the initiator session identifier. This option is required if more than one iSCSI initiator system is communicating with the same target.

10. Select **CHAP** to use CHAP for encryption or **None** from the **Authentication Type** option.
11. If the authentication type is **CHAP**, select **One Way** or **Mutual** from the **CHAP Type** option.
12. If CHAP is one way, specify **CHAP Name** and **CHAP Secret**. If CHAP is mutual, you must specify **Reverse CHAP Name** and **Reverse CHAP Secret** as well.

iSCSI Boot Workflow

The following flow chart describes the iSCSI boot work flow. After the **iSCSI Device1** is enabled and either of the connections is configured according to the above procedure, BIOS iSCSI boot operates as follows:

1. BIOS verifies if the **iSCSI initiator name** is present and that the format of the initiator name is accurate. If the initiator name is present and accurate, proceed. Otherwise, iSCSI boot exits and displays the following error message:

```
Unable to perform iSCSI boot because the iSCSI initiator name is invalid.
```
2. Provide parameters for **Interface, Protocol, VLAN, VLAN ID, VLAN Priority, Retry Count, and Timeout**.
3. Check if the **DHCP** option is enabled. If the option is disabled, proceed to the next step. If the option is enabled, BIOS tries to obtain initiator configuration information from the DHCPv4 or DHCPv6 server depending on the **Protocol** parameter obtained in the previous step. BIOS then performs the following additional steps:
 - Verifies if the connection of DHCP server to the initiator system is present. If the connection is present, proceed; otherwise, BIOS exits and displays the following error message:

```
Unable to perform iSCSI boot on connection 1/2 because the network connection to DHCPv4 (DHCPv6) server is not established.
```
 - Determine if **Target info via DHCP** option is enabled. If the option is disabled, proceed to the next step. If the option is enabled, BIOS tries to obtain target configuration information from the DHCP server. The target information obtained from the DHCP server is then verified. If the information is accurate, proceed; otherwise, it exits and displays the following error message:

```
Unable to perform iSCSI boot on connection 1/2 because the iSCSI target configuration data obtained from DHCPv4 (DHCPv6) server is invalid.
```
 - Attempts to obtain initiator information from the DHCP server. If the attempt is successful, proceed; otherwise, it exits and displays the following error message:

```
Unable to perform iSCSI boot on connection 1/2 because the BIOS cannot obtain the initiator information from DHCPv4 (DHCPv6) server.
```
4. Verify whether the network cable of the server system is present. If the network cable is present, proceed; otherwise, it exits and displays the following error message:

```
Unable to perform iSCSI boot on connection 1/2 because the network cable is not connected.
```

5. Check whether **CHAP** is enabled. If CHAP is disabled, proceed; otherwise, check whether **CHAP Type** is **One Way** or **Mutual**. For one way CHAP, obtain **CHAP Name** and **CHAP Secret** and then proceed; for mutual CHAP, obtain **Reverse CHAP name** and **Reverse CHAP secret** as well, then proceed to the next step.
6. Create a TCP connection to the iSCSI target. If a TCP connection is established within the defined retry count, proceed; otherwise, it exits and displays the following error message:

```
Unable to perform iSCSI boot on connection 1/2 because the TCP connection timed out when connecting to the iSCSI target.
```
7. Perform iSCSI login to the target system. If login is successful, proceed with iSCSI boot; If login is unsuccessful, it exits and displays the following error message:

```
Unable to perform iSCSI boot on connection 1/2 because the iSCSI login operation did not complete successfully.
```
8. Perform iSCSI boot. Load the executable image to the memory of the server and when the operation is complete, start the image.

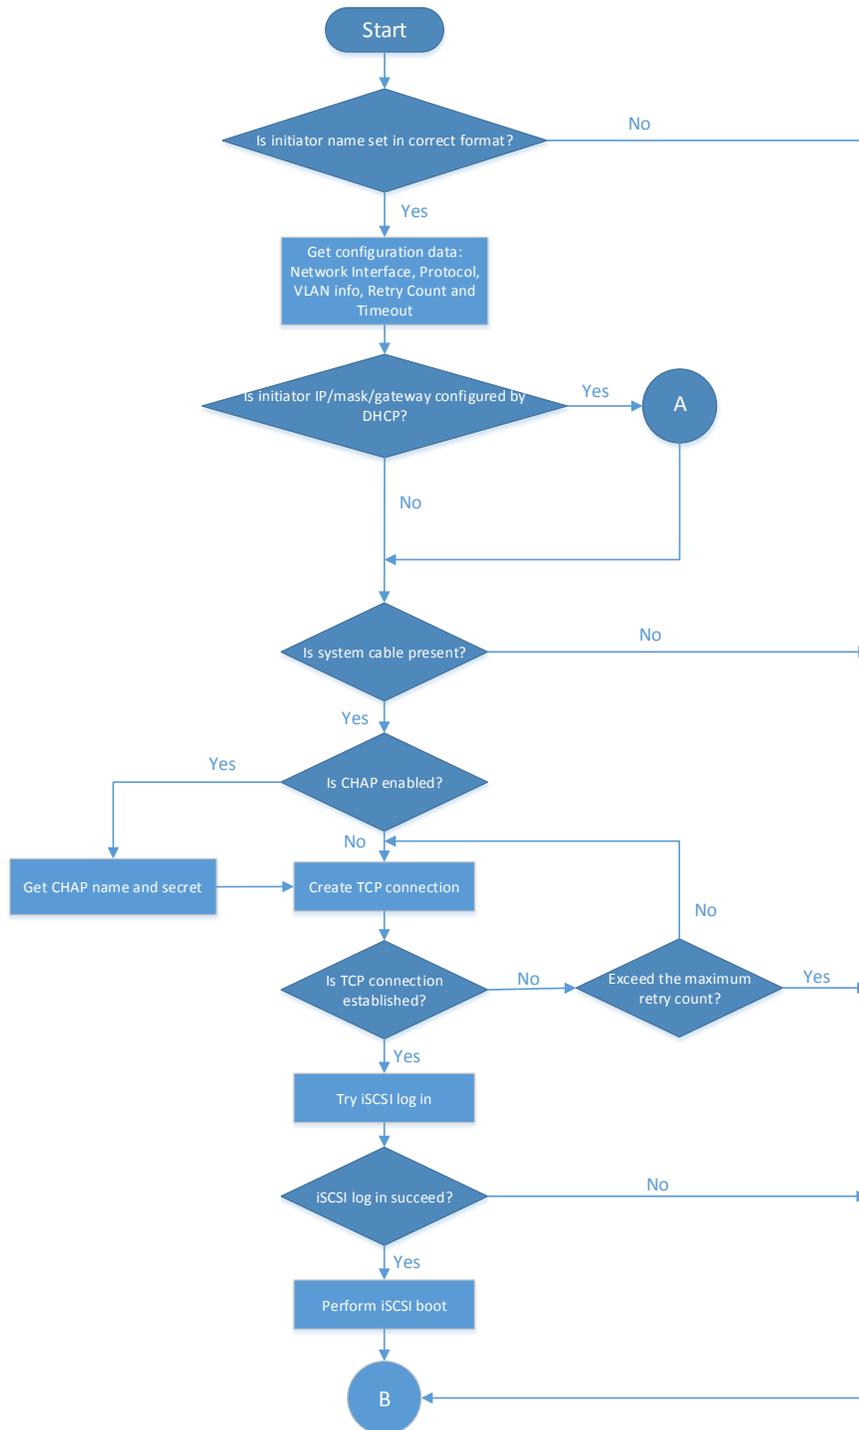


Figure 7. iSCSI boot work flowchart (1)

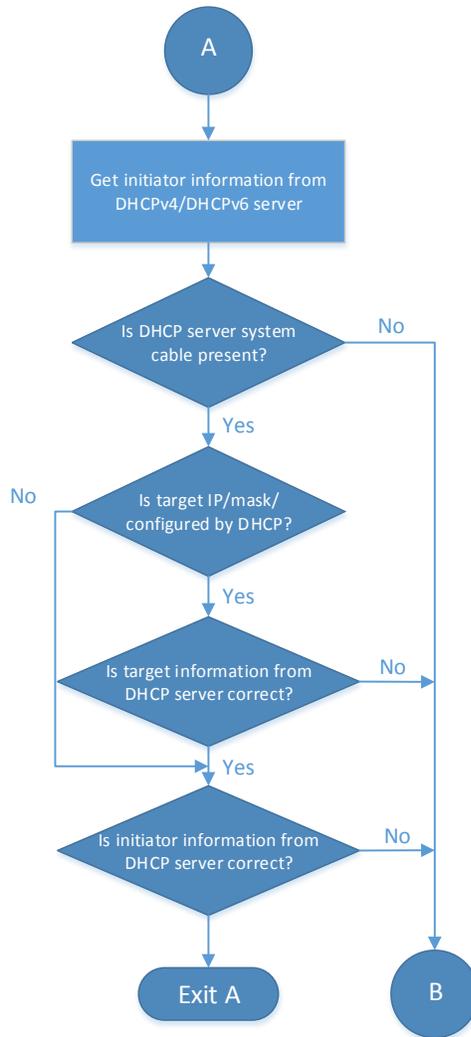


Figure 8. iSCSI boot work flowchart (2)

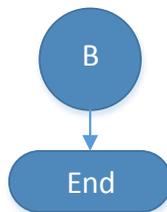


Figure 9. iSCSI boot work flowchart (3)

Current Limitations

At present, the PowerEdge iSCSI software initiator does not support redirection to a new IP address by the iSCSI target while using IPv6.

Under some circumstances such as during load-balancing, the iSCSI target will attempt to redirect an initiator to another target system using a new IP address. In such cases, PowerEdge iSCSI initiators using IPv6 will fail to reestablish a TCP connection and the iSCSI boot will fail.

Conclusion

The new iSCSI boot feature of the 13th generation of PowerEdge server BIOS enables device and vendor independent iSCSI boot by removing the limitations of vendor-provided option ROM iSCSI boot stacks. With the growing importance of network boot to support server deployment, configuration, and the bootstrap of operating environments for compute grids and software-defined services, the 13th generation of PowerEdge server iSCSI boot increases the choice of network interfaces while simplifying IT operations.

Dell continues to provide our customers with products that simplify and streamline their IT processes, reduce time and effort of the administrator to focus on activities that help to improve the organization. For more information on PowerEdge servers and IT automation, see [Dell TechCenter](#).

Appendix - Acronyms and Terminology

BIOS: Basic Input / Output System. The BIOS contains programs that are stored on a flash memory chip. The BIOS initializes the systems and loads the operating system.

Legacy BIOS: The de facto 16-bit interface between the OS and the platform firmware. It is based on the original PC BIOS.

UEFI: Unified Extensible Firmware Interface. An industry standard specification that defines a software interface between an operating system and platform firmware. UEFI is replacing Legacy BIOS as the software interface to the OS.

LOM: LAN on Motherboard. Network interfaces that are integrated directly into a motherboard of the system, as opposed to a plug-in NIC.

NDC: Network Daughter Card. Network interfaces that exist on a daughter card as opposed to being integrated directly into the motherboard or added through a plug-in NIC. Many PowerEdge systems use an NDC instead of integrating the LAN support into the motherboard (see LOM).

NIC: Network Interface Card. Typically a plug-in PCIe network adapter, but may also refer to NDC or LOM's.

PXE: Preboot eXecution Environment. A specification that allows booting from the network by identifying and then downloading the correct boot image for a system.

iSCSI: A protocol to carry SCSI disk commands over TCP/IP.

IQN format: iSCSI Qualified Name, which contains the *"iqn."* string, a date code in *yyyy-mm* format, a period ("*."*"), the reversed domain name of the naming authority, and an optional string within the character set and length boundaries prefixed with a colon ("*:"*").

VLAN: Virtual Local Area Network. A technology that uses IEEE 802.1Q tagging to identify packets on a single physical connection as belonging to separate virtual networks.

LUN: Logical Unit Number. A target system can contain multiple Logical Units, and address of each Logical Unit within a target is called a LUN.

TCP: Transmission Control Protocol. A highly reliable host-to-host protocol.

DHCP: Dynamic Host Configuration Protocol. It provides configuration parameters to an internet host.

CHAP: Challenge-Handshake Authentication Protocol. It is used to periodically verify the identity of the peer using a 3-way handshake.