

OPEN

Compute Project

**Proposed Charter
For
Open Network Install Environment**

Draft: July 2013

Version: 0.2

1 Revision History

Date	Name	Description
May 2013	Curt Brune	Initial Draft
June 25, 2013	Curt Brune	Significant updates after first round of internal review
July 2, 2013	Curt Brune	Updates for Installer and Hardware Standards

2 Contents

1	Revision History.....	2
2	Contents	3
3	Overview	4
3.1	License	4
4	Mission Statement	5
5	Design Approach	6
5.1	Overview.....	6
5.2	Initial System Boot.....	7
5.3	Subsequent System Boots	8
5.4	Locating an Installer Image	8
5.5	Installer Execution Environment.....	9
5.6	Installer Overview	10
5.7	Network Operating System Interface	10
5.8	Updating ONIE	10
5.9	Hardware Standards	10
6	Roadmap	11
7	Community Organization	12



3 Overview

The Open Network Install Environment (ONIE) is an open source initiative that defines an open “install environment” for bare metal network switches, such as the upcoming OCP Network Switch design. ONIE enables a bare metal network switch ecosystem where end users have a choice among different network operating systems.

Traditionally Ethernet switches are procured with pre-installed, captive operating systems, effectively creating networking appliances that lock end-users into a vertical supply chain. While there is industry wide discussion of a mythical "white-box" network switch, this ecosystem does not exist; failing in part due to the fact that the hardware cannot accept multiple operating systems.

As a matter of form, these switches have a management subsystem, based on a variety of CPU architectures that typically include serial console, out-of-band Ethernet, and sometimes even pluggable mass storage. This subsystem can function independently from the switching ASIC(s) associated with the "front-panel" Ethernet interfaces (i.e. without a full networking OS).

ONIE defines an open source "install environment", that runs on this management subsystem utilizing facilities in a Linux/BusyBox environment. This environment allows end-users and channel partners to install the target network OS as part of data center provisioning, in the fashion that servers are provisioned.

ONIE enables switch hardware suppliers, distributors and resellers to manage their operations based on a small number of hardware SKUs. This in turn creates economies of scale in manufacturing, distribution, stocking, and RMA enabling a thriving ecosystem of both network hardware and operating system alternatives.

Highlights:

- Combines a boot loader with a modern Linux kernel and BusyBox
- Provides an environment for installing any network OS
- Disruptive, liberating users from a captive, pre-installed network OS
- Helps automate large scale data center switch provisioning
- Manage your switches like you manage your Linux servers

3.1 License

ONIE is licensed under the GPLv2 license.

4 Mission Statement

The mission of ONIE is to:

“Provide an open installation environment where multiple network operating system vendors can execute a compliant network OS installer. The environment facilitates process automation, scalability and ease of provisioning.”

Fulfilling this mission will help enable an ecosystem where end users can mix and match switching hardware from multiple hardware vendors with a variety of network operating systems. For any particular hardware platform ONIE provides an environment where network OS vendors can execute a compliant network OS installer.

Expanding on the mission statement, other high level goals include:

- Easy for hardware vendors to provide ONIE pre-installed
- Easy for software vendors to provide ONIE enabled OS installers
- Easy for end users to use and configure out of the box
- Flexibility and extensibility for advanced end users
- Process automation and scalability
- Leverage existing open standards whenever possible
- Coordinate with other groups within the OCP to drive efforts

5 Design Approach

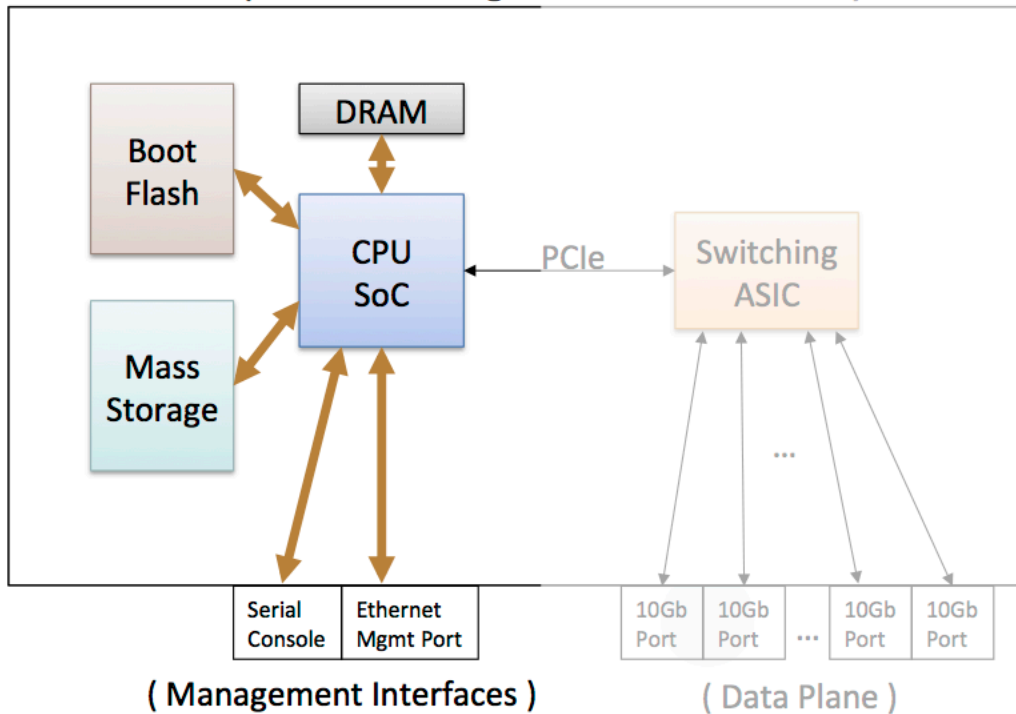
Using the Mission Statement as a guide the approach to delivering and maintaining ONIE includes:

- Any solution must scale to 1,000+ switches
- Use existing standards whenever possible
- Provide a 100% open reference implementation of ONIE for an open hardware platform
- Encourage and work with switch hardware vendors and network operating system vendors to promote the use of ONIE
- For every feature or function consider the role of security and ensure system integrity
- Contribute to the definition of features, attributes and interface standards for OCP hardware designs
- Work with OCP hardware designers to implement required functionality
- As part of the software deliverable include a methodology to validate the functionality
- Provide documentation, including specifications, tutorials, user guides and FAQs

5.1 Overview

ONIE is the combination of a boot loader and a small operating system for bare metal network switches that provides an environment for automated provisioning. ONIE utilizes the CPU complex of the switch, but not the forwarding data plane as shown in the following figure:

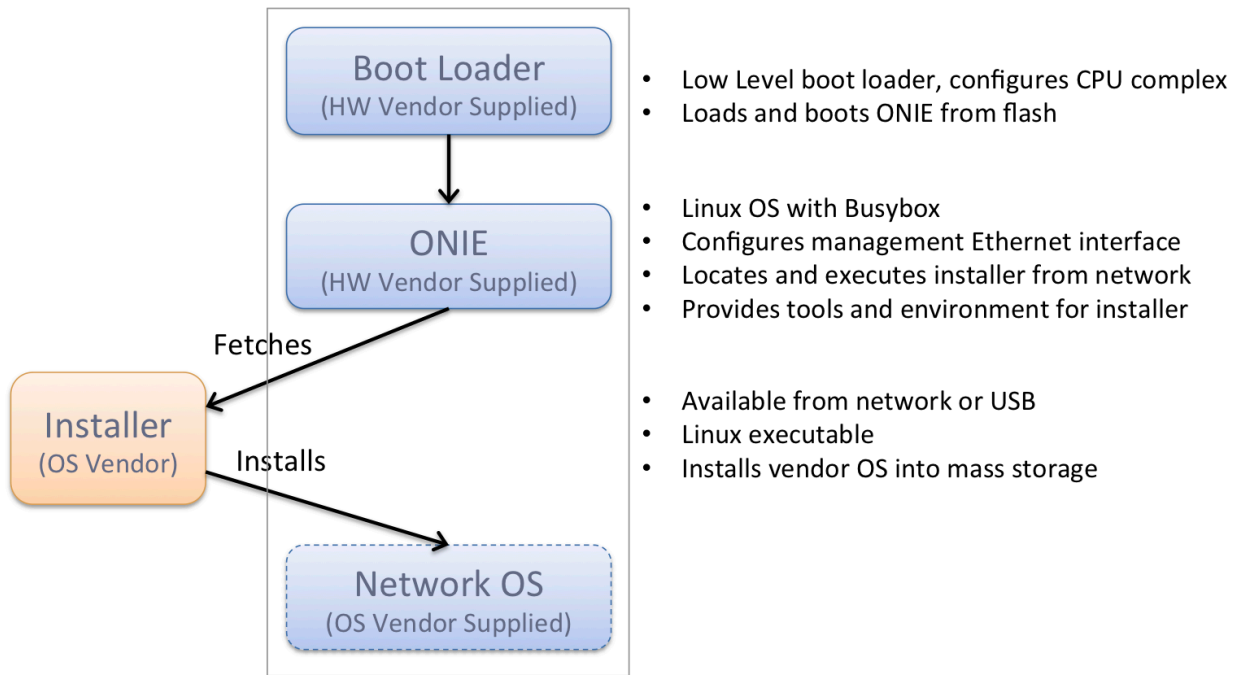
CPU Root Complex and Management Interfaces Only



5.2 Initial System Boot

When a new machine boots for the first time ONIE locates and executes a network OS vendor's installation program, as shown here:

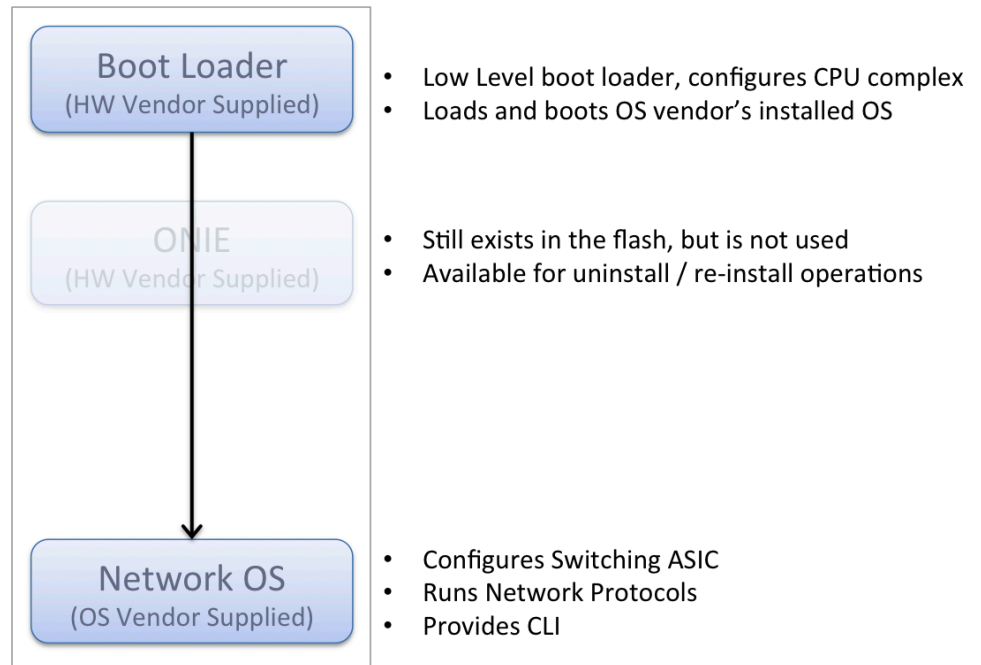
First Time Boot Up



5.3 Subsequent System Boots

ONIE is **not** used on every boot of the system. After initial installation subsequent boots go straight into the network OS, bypassing ONIE. This is depicted here:

Vendor's OS is Already Installed. ONIE is **not** used.



Mechanisms exist for a system to re-enter the installation phase. An API will be defined so that network operating systems can direct the system to re-enter the installation phase.

5.4 Locating an Installer Image

ONIE uses a number of methods to locate a network OS installer, including, but not limited to:

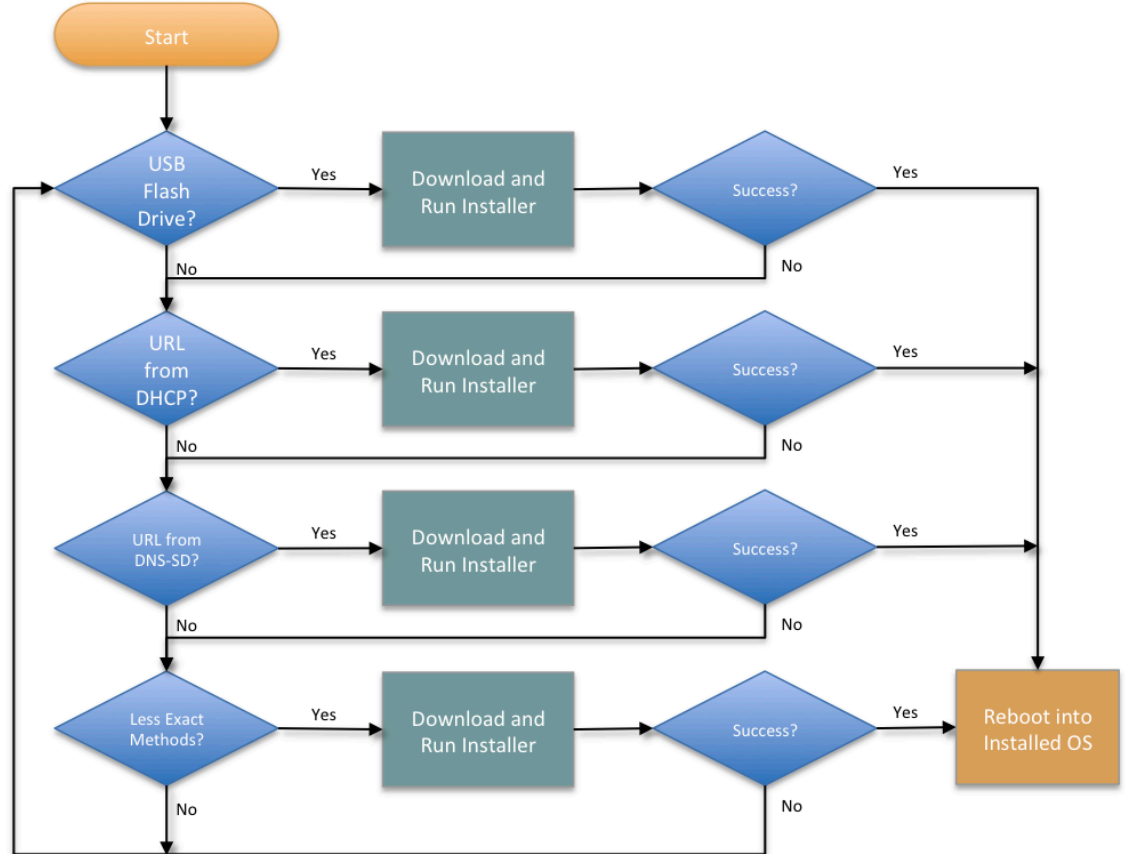
- Statically configured from the boot loader
- Locally attached storage, for example a USB memory stick
- DHCPv4 / DHCPv6
- IPv4 / IPv6 link local neighbors
- mDNS / DNS-SD
- PXE-like TFTP waterfall

The preferred method for image download is HTTP as it offers robust performance for large image sizes. TFTP is also supported, but its use is discouraged in favor of HTTP.

Once an image is located ONIE proceeds to execute the network OS installer.

The following diagram illustrates an example of the location and execution steps:

Discovery Methods



In the previous diagram the “Less Exact Methods” box refers to mechanisms that use probing techniques to locate an image, such as:

- IPv4 / IPv6 link local neighbors
- PXE-like TFTP waterfall

5.5 Installer Execution Environment

When executing the installer, ONIE exports a number of environment variables for the installer to use. These variables include system identification information as well as information learned during the image discovery process.

An example of the information exported to the installer includes:

- Installer URL
- HW vendor identifier
- HW platform identifier
- HW serial number
- Ethernet management MAC address
- IP address (from DHCP)
- Hostname (from DHCP)

These variables allow an installer to integrate with other process automation and orchestration, immediately tying together serial numbers, MAC addresses and network OS versions.

5.6 Installer Overview

The installer's responsibility is to persistently install the operating system into the hardware. In fulfilling that role an ONIE compliant network OS installer has a lot of flexibility.

Some examples of what an installer could do:

- Fetch additional binaries and configuration files via HTTP
- Chat with inventory control systems via HTTP
- Download a new kernel+initramfs and kexec(8) into it

5.7 Network Operating System Interface

ONIE provides an environment for a network OS to perform system provisioning and maintenance operations, including:

- Reinstalling a different network OS
- Uninstalling, i.e. wiping the system clean
- Rescue and recovery
- Updating ONIE

This environment forms an interface between ONIE and the network OS.

5.8 Updating ONIE

ONIE provides a mechanism for updating itself. This mechanism proceeds much like the network installer discovery and executing phase, however the image in this case is an *ONIE update image*. Once located, the ONIE update image is downloaded and executed.

5.9 Hardware Standards

In order to form a solid foundation for an open hardware switch platform, various standards need to be developed and ratified. As a component of the open hardware switch, ONIE, would contribute to and push standards that define the hardware/software interface.

As an example, an open board information EEPROM format would greatly benefit the bare metal switch community. The EEPROM is programmed during manufacturing by the hardware vendor and contains information such as:

- Serial Number
- Product Identification
- MAC address allocation

This information is needed by both ONIE and network OS vendors. Having a common format simplifies operations for hardware vendors, ONIE implementers and network OS vendors.

6 Roadmap

The future of ONIE follows a number of concurrent paths:

- Continue to drive the adoption of ONIE for currently shipping hardware products and network operating systems
- Adapt ONIE to new OCP hardware platforms, including new CPU architectures
- Enhance and amend ONIE in response to OCP community feedback
- Continue to improve the documentation for developers and users

7 Community Organization

ONIE development will utilize a number of strategies borrowed from other successful open source software projects, including:

- Provide a central, public website for the project
- Provide documentation, tutorials, user guides and FAQs
- Provide an open membership mailing list for discussions, with an archive
- Provide a moderated issue tracking system
- Provide a public git repository for the core ONIE source code. Anyone can read from the repository, but only key project members will have commit permissions