

Bare-metal provisioning in data center, including network switches

Steven Shiau, Chen-Kai Sun, Yu-Chin Tsai,
NCHC, Taiwan

Jim Hao Chen, Fei Yeh, Joe Mambretti
ICAIR, Northwestern University, USA

Also at: OCP Global Summit
2019, March 14, San Jose
Clonezilla for Open Network
Install Environment (ONIE):
A Network Switch Storage
Operations Enabler
Presenter: Jim Chen

steven@nchc.org.tw; jim-chen@northwestern.edu

March 15, 2019




System Deployment/Provisioning



image source: maggiesfarm.anotherdotcom.com
www.compsults.com, and jervisabreo.com

What is Clonezilla?

- A partition and disk imaging/cloning utility similar to True image® or Ghost®
- GPL license
- A system provisioning, bare metal recovery tool for



*1



*2



*3




*4

VMFS

VMware
ESX/ESXi

*5




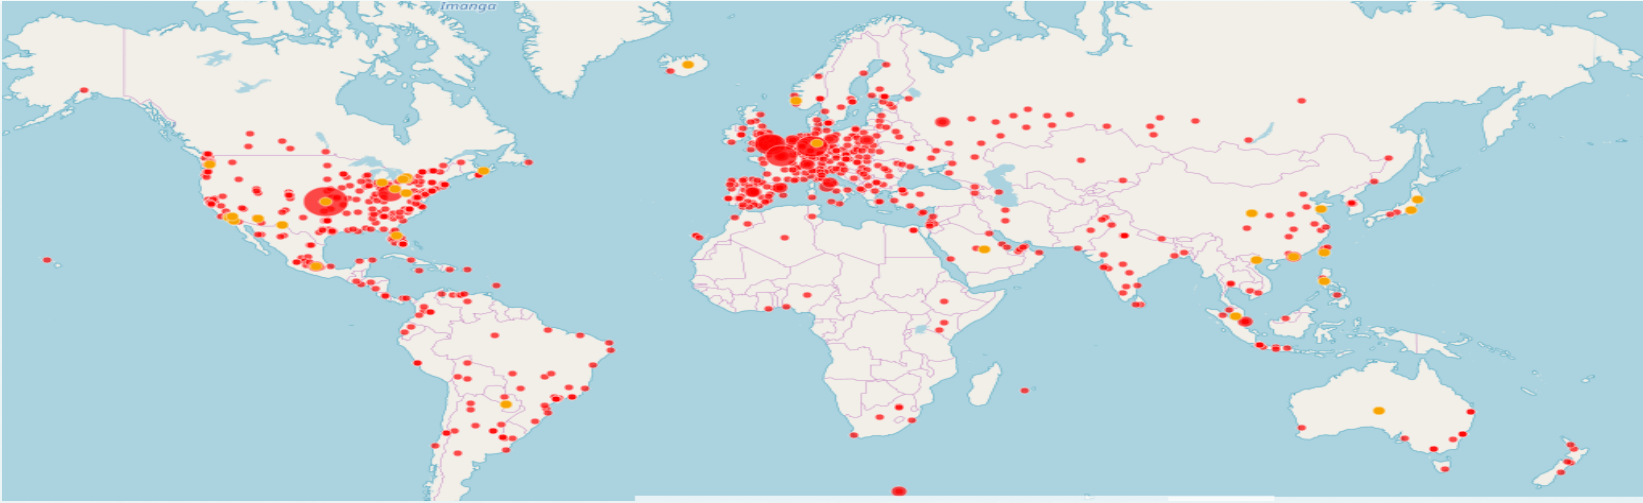
MINIX

*6



*Logo source: (1) Larry Ewing, Simon Budig and Anja Gerwinski, (2) Apple, (3) Microsoft, (4) Marshall Kirk McKusick, (5) VMWare (6) Distrowatch.com

Clonezilla



**19 million downloads
(2004-2018)**



Machine Room Deployment & Provisioning




- Server nodes
- Compute nodes
- Storage nodes



- Network switches



until...

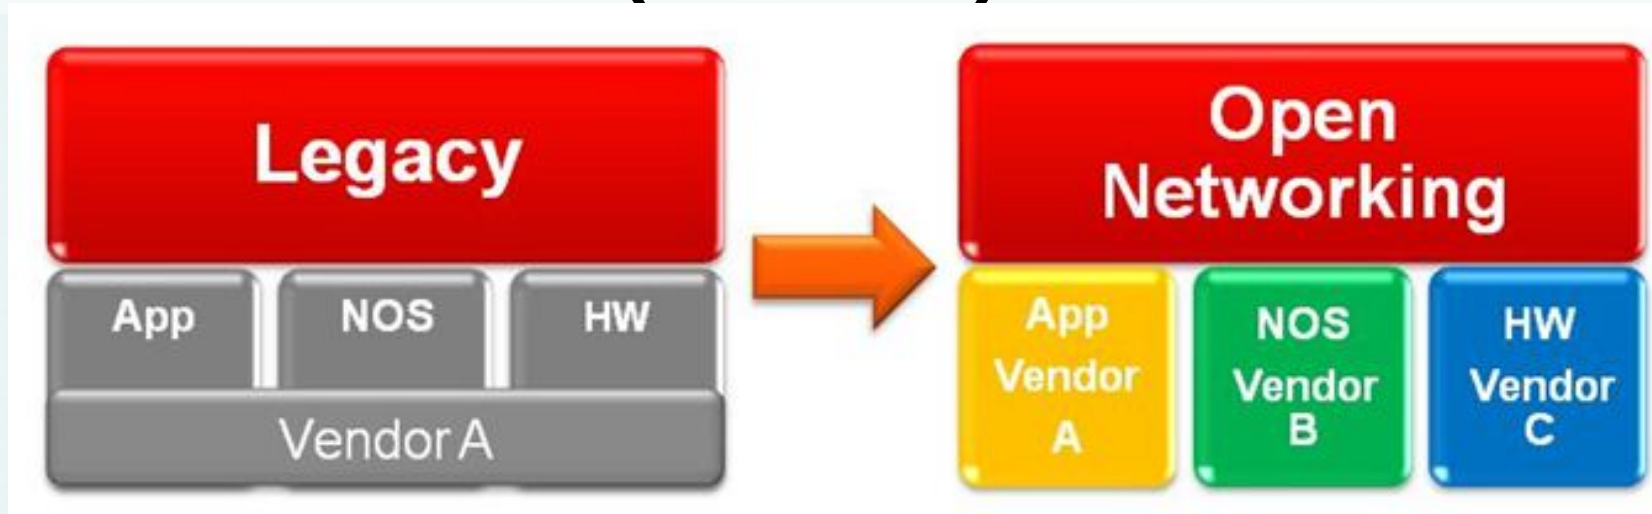


Clonezilla

onie



Open Network Install Environment (ONIE)



Created by
Cumulus Networks
in 2012

Adopted by
Open Compute
project in 2013

bare metal,
white box

open source,
open networking
hardware

onie

ONIE compatible NOS

| Vendor | OS | Base |
|---------------------|----------------------------------|----------------|
| Big Switch Networks | Switch Light | Debian |
| Cumulus Networks | Cumulus Linux | Debian |
| Dell | Force 10 Operating System (FTOS) | NetBSD |
| HPE | OpenSwitch | Yocto |
| Mellanox | MLNX_OS | Debian |
| Lenovo | cNOS | Yocto |
| Lenovo | eNOS | Yocto |
| OCP | ONL | Debian |
| Pica8 | PicOS | Debian |
| Pluribus Networks | Netvisor OS | Illumos/CentOS |

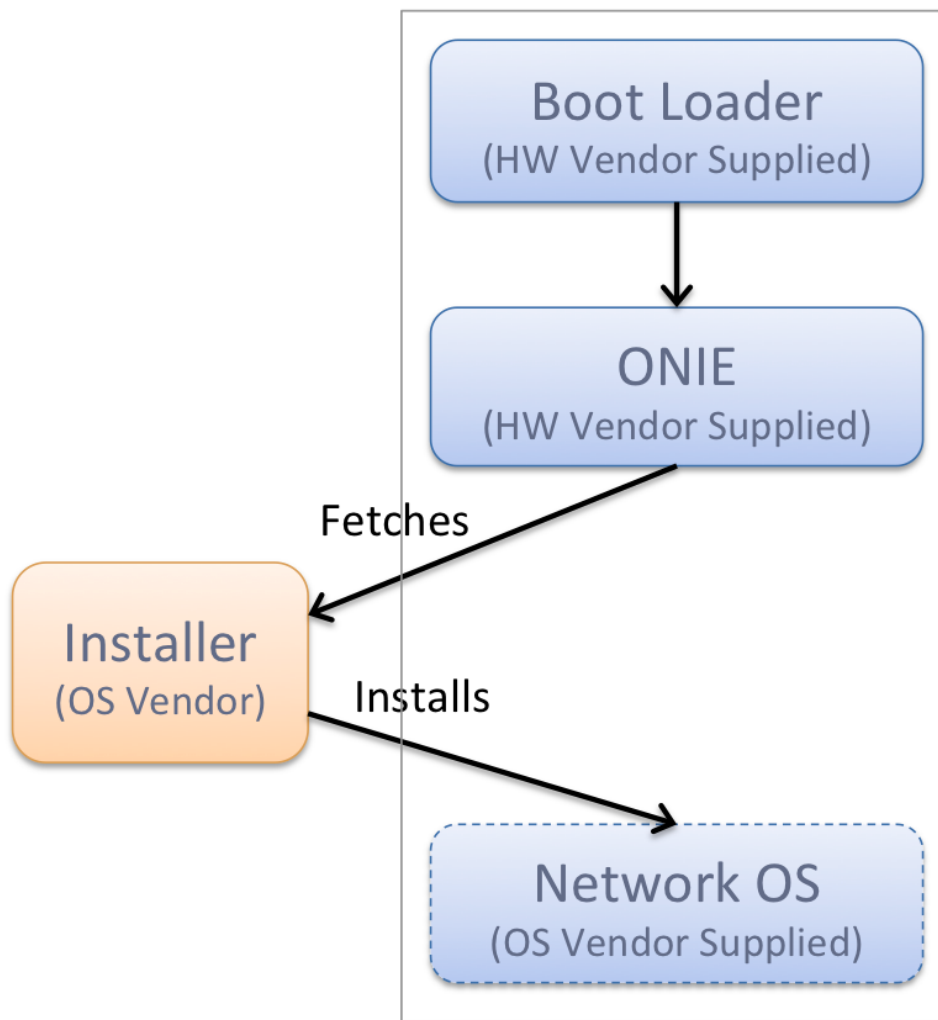
CPU (32/64-bit)

- ✓ X86
- ✓ PowerPC
- ✓ ARM



Inside ONIE: initial system boot

First Time Boot Up



- Low Level boot loader, configures CPU complex
- Loads and boots ONIE from flash
- Linux OS with Busybox
- Configures management Ethernet interface
- Locates and executes installer from network
- Provides tools and environment for installer
- Available from network or USB
- Linux executable
- Installs vendor OS into mass storage

NOS Interface


where Clonezilla live starts...

| | |
|--------------|---|
| Reinstalling | Reinstalling a different NOS |
| Uninstalling | Uninstalling/wiping the system clean |
| Rescue | Rescue and recovery |
| Updating | Updating ONIE |
| | Updating the machine CPLD/FPGAs and BIOS firmware |



Inside ONIE: subsequent system boots

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



- Low Level boot loader, configures CPU complex
- Loads and boots OS vendor's installed OS

- Still exists in the flash, but is not used
- Available for uninstall / re-install operations

- Configures Switching ASIC
- Runs Network Protocols
- Provides CLI

Prepare Clonezilla live for ONIE

Once per software version

Download Clonezilla
live ISO

Convert to ONIE self-
extracting boot file

Clonezilla live image

NOS Image store
(Networked or locally
attached)

ONIE network switch

Boot loader

ONIE


Clonezilla
live

Target NOS

Fetch (http, ftp, USB...)

Copy to Memory

Store, Recover,
Provision req.



Screenhots from Dell S4048-ON ONIE switch



```
GNU GRUB version 2.02-beta2+e4a1fe391

+-----+
| Open Network Linux |
| *ONIE               |
+-----+

Use the ^ and v keys to select which entry is highlighted.
Press enter to boot the selected OS, 'e' to edit the commands
before booting or 'c' for a command-line.
```

```
GNU GRUB version 2.02-beta2+e4a1fe391

+-----+
| ONIE: Install OS   |
| *ONIE: Rescue      |
| ONIE: Uninstall OS |
| ONIE: Update ONIE  |
| ONIE: Embed ONIE   |
| PLATFORM-DELL-DIAG |
+-----+

Use the ^ and v keys to select which entry is highlighted.
Press enter to boot the selected OS, 'e' to edit the commands
before booting or 'c' for a command-line.
```

```
Build Date: 2016-09-07/15:54-0700
Info: Mounting kernel filesystems... done.
Info: Mounting ONIE-BOOT on /mnt/onie-boot ...
Info: Using eth0 MAC address: e4:f0:04:91:37:a7
Info: Using eth1 MAC address: e4:f0:04:91:37:a8
Info: eth0: Checking link... up.
Info: Trying DHCPv4 on interface: eth0
ONIE: Using DHCPv4 addr: eth0: 192.168.0.94 / 255.255.255.0
Info: eth1: Checking link... down.
ONIE: eth1: link down. Skipping configuration.
ONIE: Failed to configure eth1 interface
Starting: dropbear ssh daemon...random: dropbear urandom read with 32 bits of entropy available
done.
Starting: telnetd... done.
discover: Rescue mode detected. Installer disabled.

Please press Enter to activate this console.
To check the install status inspect /var/log/onie.log.
Try this: tail -f /var/log/onie.log

** Rescue Mode Enabled **
ONIE:/ # random: nonblocking pool is initialized
```

```
ONIE:/ # onie-nos-install http://stevenshiau.org/clonezilla-live/onie/clonezilla-live-2.6.0-37-amd64.iso-onie.bin
discover: Rescue mode detected. No discover stopped.
Info: Fetching http://stevenshiau.org/clonezilla-live/onie/clonezilla-live-2.6.0-37-amd64.iso-onie.bin ...
Connecting to Stevenshiau.org (140.110.17.130:80)
installer 100% |*****| 273M 0:00:00 ETA
ONIE: Executing installer: http://stevenshiau.org/clonezilla-live/onie/clonezilla-live-2.6.0-37-amd64.iso-onie.bin
Verifying image checksum ... OK.
Preparing image archive ... OK.
auto-detecting console...
using console=ttyS0,115200n8
Loading new kernel ...
kexec --load --initrd=initrd.img --append="boot=live nopersistent noeject dhcp fromiso=/conf/live.iso live-media=ram ocs_live_run=ocs-live-general live-getty console=ttyS0,115200n8" vmlinuz
sd 4:0:0:0: [sda] Synchronizing SCSI cache
kexec: Starting new kernel
[ 0.000000] Linux version 4.19.0-1-amd64 (debian-kernel@lists.debian.org) (gcc version 8.2.0 (Debian 8.2.0-13)) #1 SMP Debian 4.19.13-1 (2018-12-30)
[ 0.000000] Command line: boot=live nopersistent noeject dhcp fromiso=/conf/live.iso live-media=ram ocs_live_run=ocs-live-general live-getty console=ttyS0,115200n8
```

Demo

- ONIE network switch backup or restore
 - Prepare Clonezilla live ONIE self-extracting boot file
 - `sudo ocs-iso-2-onie clonezilla-live-2.6.0-37-amd64.iso`
 - ONIE network switch
 - Simulate: Boot VM in Virtualbox with ONIE iso
 - Use serial console
 - Enter rescue mode
 - Switch to Clonezilla live
 - Backup, restore...



onie



```
› git clone https://github.com/opencomputeproject/onie.git
› cd onie/build-config
› sudo apt-get update
› sudo apt-get install build-essential
› make debian-prepare-build-host
› make -j4 MACHINE=kvm_x86_64 all recovery-iso
  ../build/images/onie-recovery-x86_64-kvm_x86_64-r0.iso
```



Reference

- <https://clonezilla.org>
- <http://onie.org>
- <https://github.com/opencomputeproject/onie>
- <https://www.opencompute.org>
- <https://www.icair.org>
- <https://www.nchc.org.tw>



Questions?

