


presented by



Microsoft



Filling UEFI/FW Gaps in the Cloud

UEFI Spring Plugfest – May 18-22, 2015
Presented by Mallik Bulusu – Microsoft
and Vincent Zimmer - Intel

Agenda

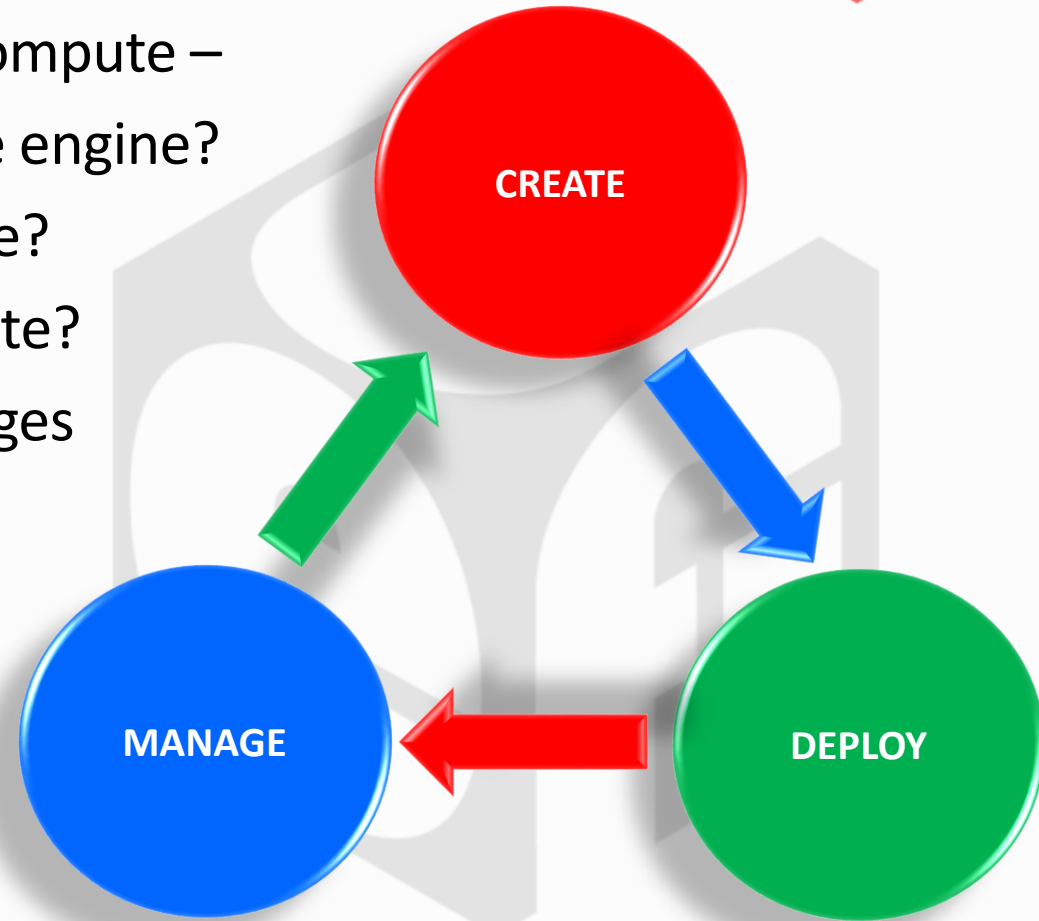


- Cloud Server Firmware Challenges
- What is Open Compute
- Intro to UEFI
- Firmware Update
- Provisioning
- Tools & Diagnostics
- Security
- Conclusion

Firmware Challenges in the Cloud



- Design constraints for the Compute –
 - How to **create** the compute engine?
 - How to **deploy** the compute?
 - How to **manage** the compute?
- Key Cloud Firmware Challenges
 - Firmware Updates
 - Bare Metal Provisioning
 - Security
 - Tools & Diagnostics



What is OCP?




Open Compute: Industry Collaboration




facebook


Open Compute: Industry Collaboration




Open Compute: Industry Collaboration




Open Compute: Industry Collaboration



Open Compute: Complex Designs



Open Compute: ~~Complex~~ Simple Designs



Open Compute: Designed for Scale



Open Compute: Why Facebook loves it



38%

energy
efficiency
gained

24%

cost
savings

Open Compute: Open Source



the Open Compute Project
<http://opencompute.org/>

Clone in Mac ZIP HTTP Git Read-Only `https://github.com/facebook/opencompute.git` Read-Only access

branch: master Files Commits Branches 1 Tags Downloads 4

Latest commit to the **master** branch

OCP 2012 Summit Motherboard presentation + ...
desertmonkey authored 10 hours ago commit 0b407686f6

[opencompute](#) /


name	age	message	history
AM		<pre>\$ git clone https://github.com/facebook/opencompute.git Cloning into opencompute...</pre>	
HV			
Int			

Open Compute and Firmware




- No recommendation for firmware
- Various solutions for boot and network provisioning
- Align on UEFI based technology

Why UEFI?



Mostly written in C.
High code re-use.




Emphasis on Specifications.
Standards compliance.




Better platform scaling. For e.g. removes shadow ROM limits.



Storage.
GPT removes 2.2 TB MBR restriction.




CPU Architecture independent. Platform design flexibility.



Secure boot solves "trust" related system integration challenges.

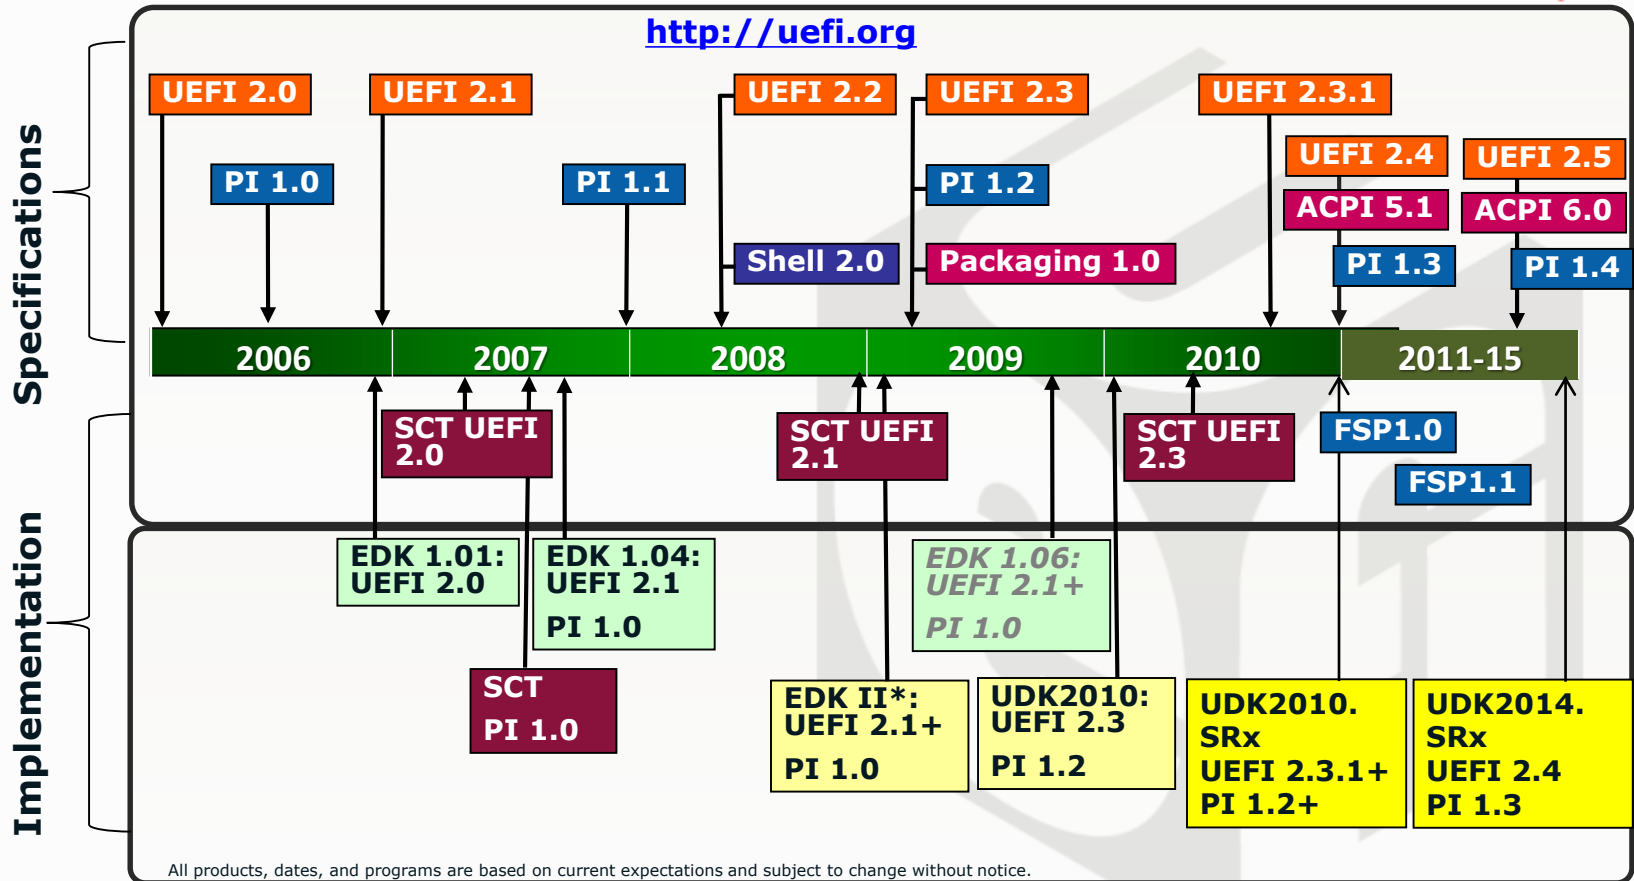


Pre-boot Networking.
Ipv4, Ipv6, PXE, VLAN, iSCSI etc.



UEFI shell improves pre-boot testing & diagnostics experience.

Timeline



<http://tianocore.org>

<https://github.com/tianocore/edk2>

How to build it



Industry Standards Compliance

- UEFI 2.0, UEFI 2.1, UEFI 2.2, UEFI 2.3, UEFI2.4; PI 1.0, PI 1.1, PI 1.2, PI1.3; ACPI 5.1

Extensible Foundation for Advanced Capabilities

- Pre-OS Security
- Rich Networking
- Manageability

Support for UEFI Packages

- Import/export modules source/binaries to many build systems

Maximize Re-use of Source Code**

- Platform Configuration Database (PCD) provides “knobs” for binaries
- ECP provides for reuse of EDK1117 (EDK I) modules
- Improved modularity, library classes and instances
- Optimize for size or speed

Multiple Development Environments and Tool Chains**

- Windows, Linux, OSX
- VS2003, VS2005, WinDDK, Intel, GCC

Fast and Flexible Build Infrastructure**

- 4X+ Build Performance Improvement (vs EDK I)
- Targeted Module Build Flexibility

** benefit of EDK II codebase

Maximize the open source at www.tianocore.org

New Specification advances for Cloud



- Error support
 - CPER
- New memory topology
 - NVDIMM
 - Reliability
- Boot from HTTP






Firmware Updates



Firmware Update Challenges



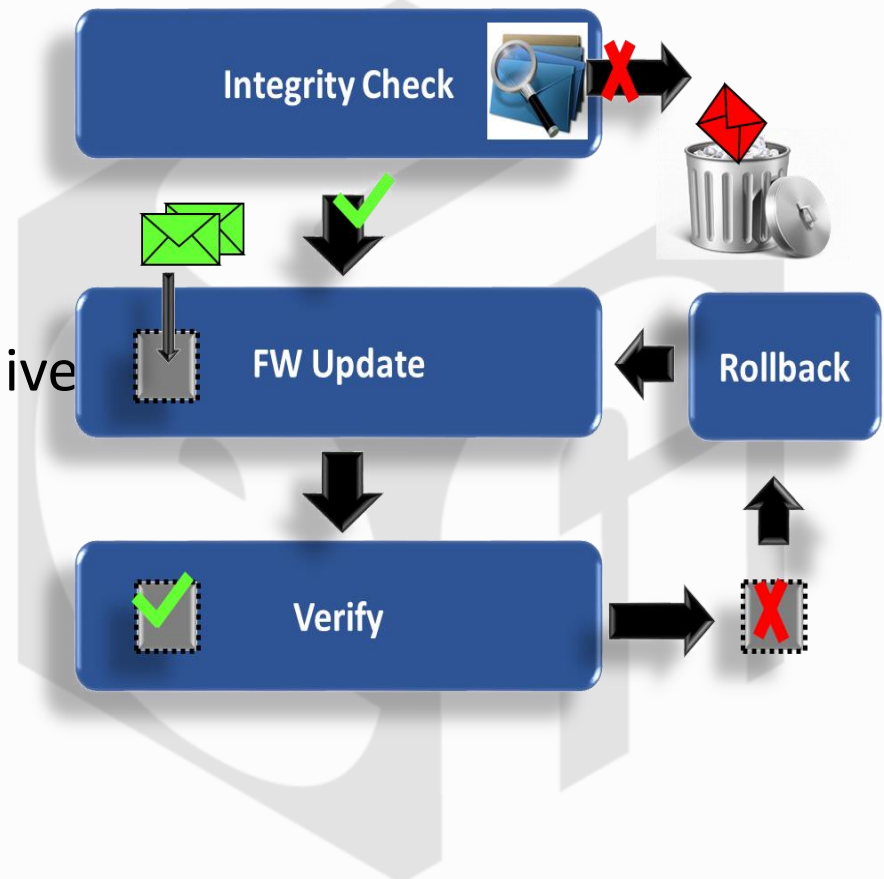
- Components from multiples vendors
- Delivering firmware
- Different types of devices
- Recovery from failures
- Node equivalence across datacenter
- Security, security, security.....



Solving the firmware update Challenge




- Reliable update story
 - Fault tolerant
 - Scalable & repeatable
- How can UEFI Help?
 - Capsule model for binary delivery
 - Bus / Device Enumeration
 - Managing updates via
 - EFI System Resource Table
 - Firmware Management Protocol
 - Capsule Signing



Delivering firmware updates



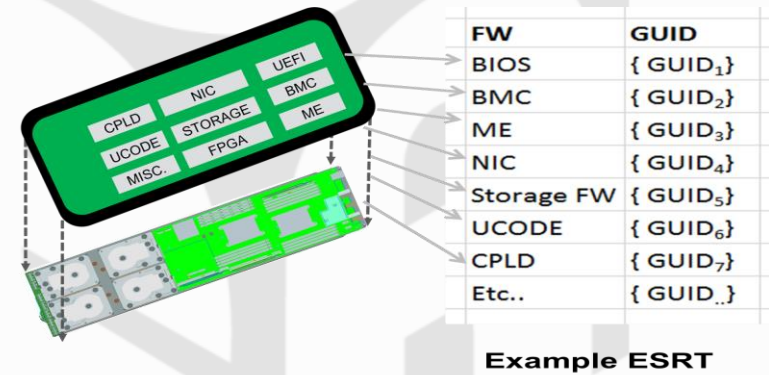
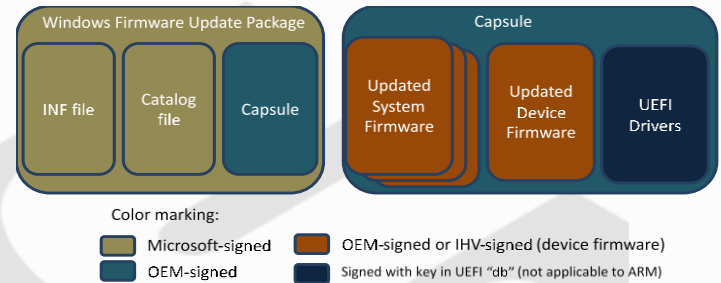
- UEFI supports Capsule format
 - Tools for capsule generation
 - Core logic for capsule handling
- Extensible Capsule format
 - Self-contained
 - Discrete updates
 - Composite updates
- Firmware Management Protocol
 - Reading / updating firmware
 - Integrity checks



EFI System Resource Table



- Update types
 - Largely OS assisted
 - Largely BIOS assisted
- FW updateability rules can be encoded into the capsule
 - Least version
 - Signing
- Describe various updateable components on the platform





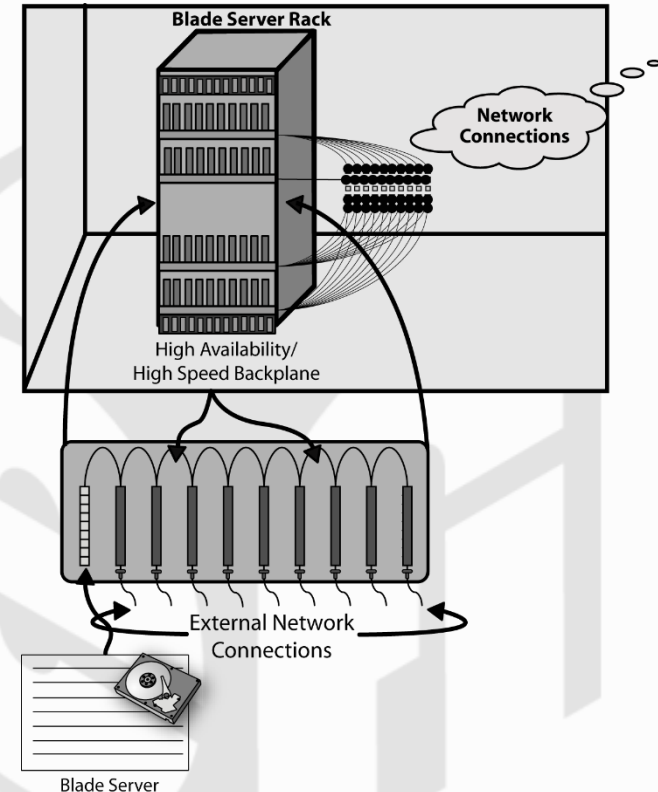
Bare Metal Provisioning



Bare Metal Provisioning Challenges



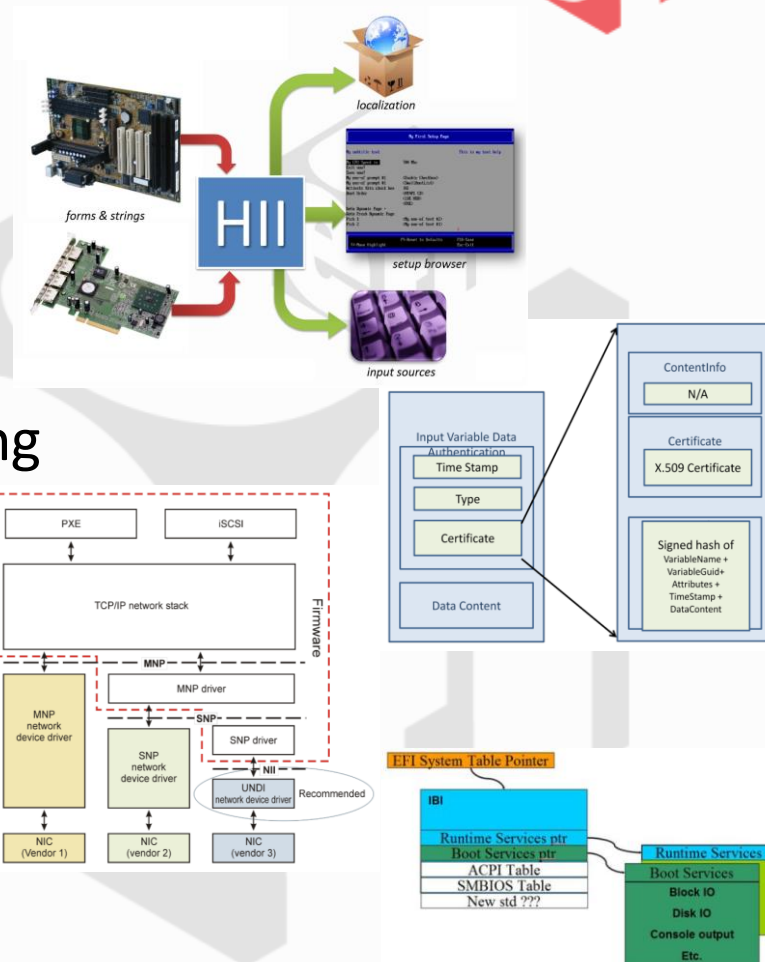
- Hardware Detection
- Installation
 - Local / Remote
- Configuration
 - Local / Remote / Scriptable
- Cloning
 - Automated
- Backup / Recovery
 - Local / Remote / Automated



Bare Metal Provisioning Solutions



- Need a 'no-touch', automated installation mechanism
 - Repurpose / Configure / Recover
- HII and IFR for consistent & scriptable configuration
- Non-blocking local disk and networking services for high throughput image delivery and recovery
- UEFI Variables for booting and Authenticated Variables for safe storage of settings, like UEFI secure boot database





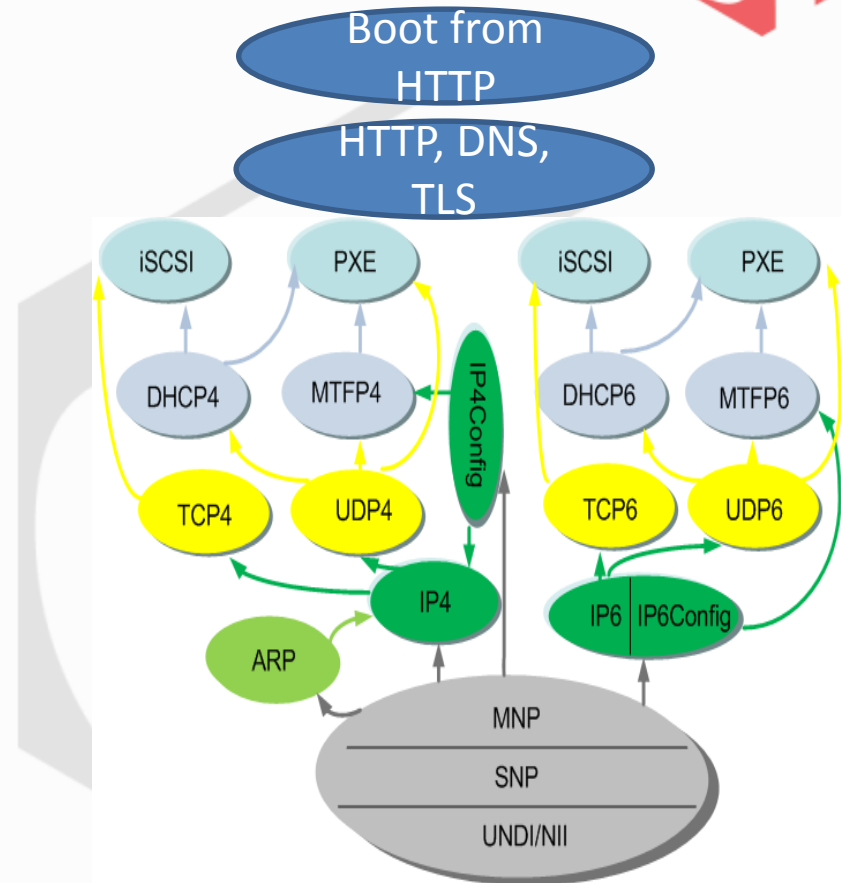
Networking in UEFI



Networking



- UEFI offers rich set of Networking Features during pre-boot
 - PXE boot support for network boot, OS installations, provisioning etc.
 - Native support for IPv4 as well as IPv6
 - Network file system support
 - Virtual LAN support, iSCSI
 - IpSec for supporting secure communication
- Evolution of networking –
 - DNS
 - TLS
 - RFC 5970 allows for 'boot from URI'
 - Boot from HTTP

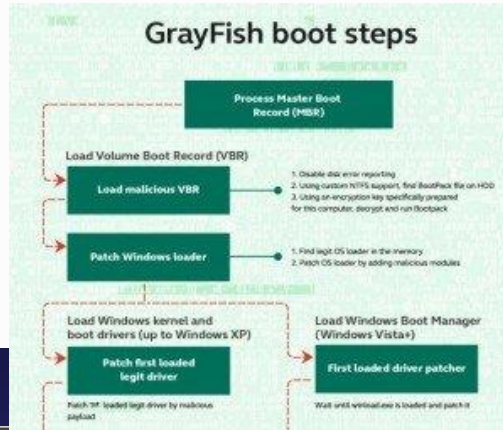




Security



Attacks on Firmware



**DE MYSTERIIS DOM JOBSIVS:
MAC EFI ROOTKITS**

SNARE
@ SYSCAN SINGAPORE
APRIL 2012

assurance

**IN CONCLUSION...
I HAD FUN.**

- ▶ So basically we're all screwed
 - ▶ What should you do?
 - ▶ Glue all your ports shut
 - ▶ Use an EFI password to prevent basic local attacks
 - ▶ Stop using computers, go back to the abacus
 - ▶ What should Apple do?
 - ▶ Implement UEFI Secure Boot (actually use the TPM)
 - ▶ Use the write-enable pin on the firmware data flash properly
 - ▶ NB: They may do this on newer machines, just not my test one
 - ▶ Audit the damn EFI code (see Heasman/TTL)
 - ▶ Sacrifice more virgins

De Mysteriis Dom Jobsivus - 5/5cm
April, 2012

**Hacking the Extensible
Firmware Interface**

John Heasman, Director of Research

Code Injection Attacks

- ▶ Important when firmware verifies digital signatures
 - Depends on implementation flaw in driver
 - e.g. stack overflow, heap overflow
 - or incorrect signature verification
- ▶ Plenty of targets:
 - File system drivers (e.g. FAT32, HFS+)
 - PE parsing code
 - Crypto code (Data in certs, ASN.1 decoding)
 - Network interaction (PXE)

Network Nightmare

Ruling the nightlife between
shutdown and boot with pxeexploit

Bootkits

Stoned Bootkit

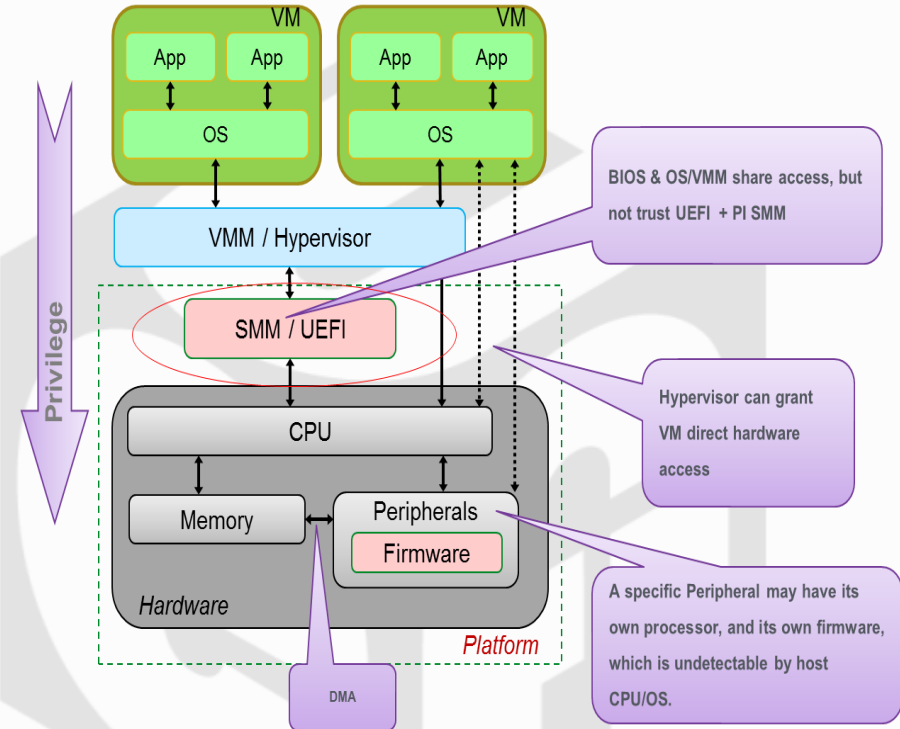
Peter Kleissner

Stephen Cobb, senior security researcher at ESET North America, says that hacking firmware can be particularly effective because it is so hard to eliminate. It's also particularly challenging to do, says Jean Taggart, security researcher at Malwarebytes. "Doing this on just one brand of hard drive would be an almost Herculean task," he says. "You have to understand the hardware as well—if not more—than the original manufacturer." — Stan Alcorn, Marketplace, Feb 17, 2015

Security



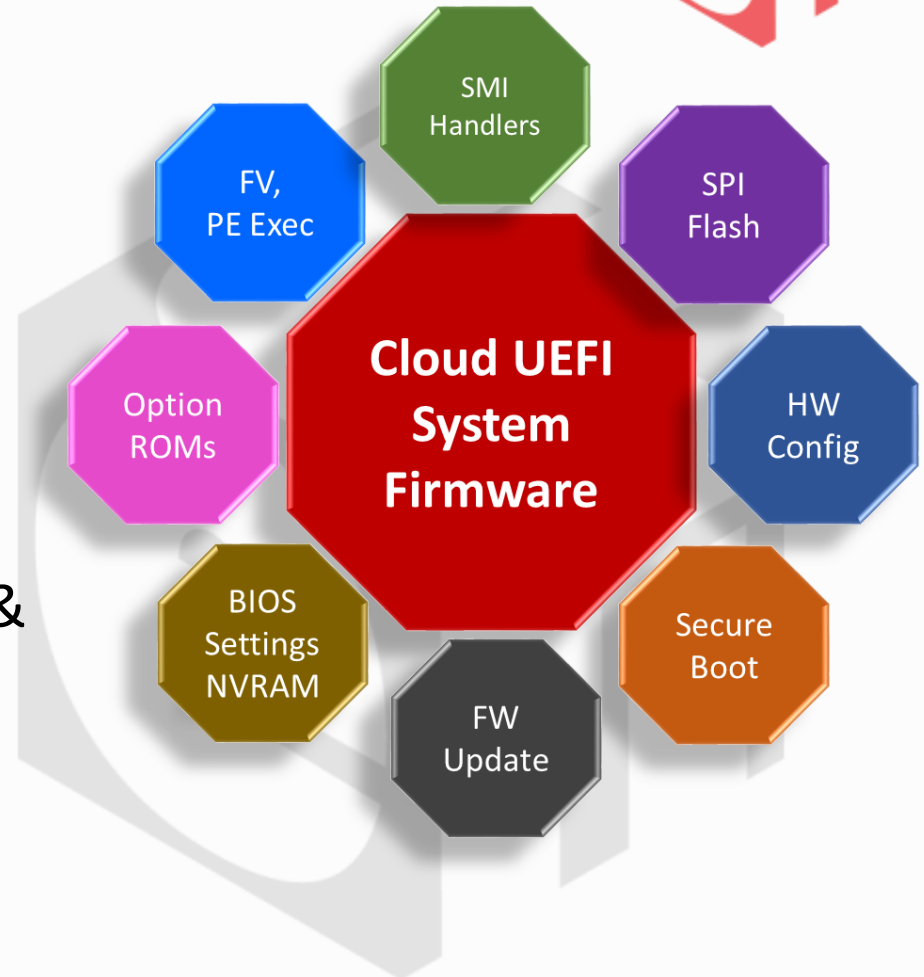
- Different elements in platform from many vendors
- How to establish trust anchor in the hardware
- How to protect elements
- How to protect the platform
- How to allow platform scaling



Security Solutions

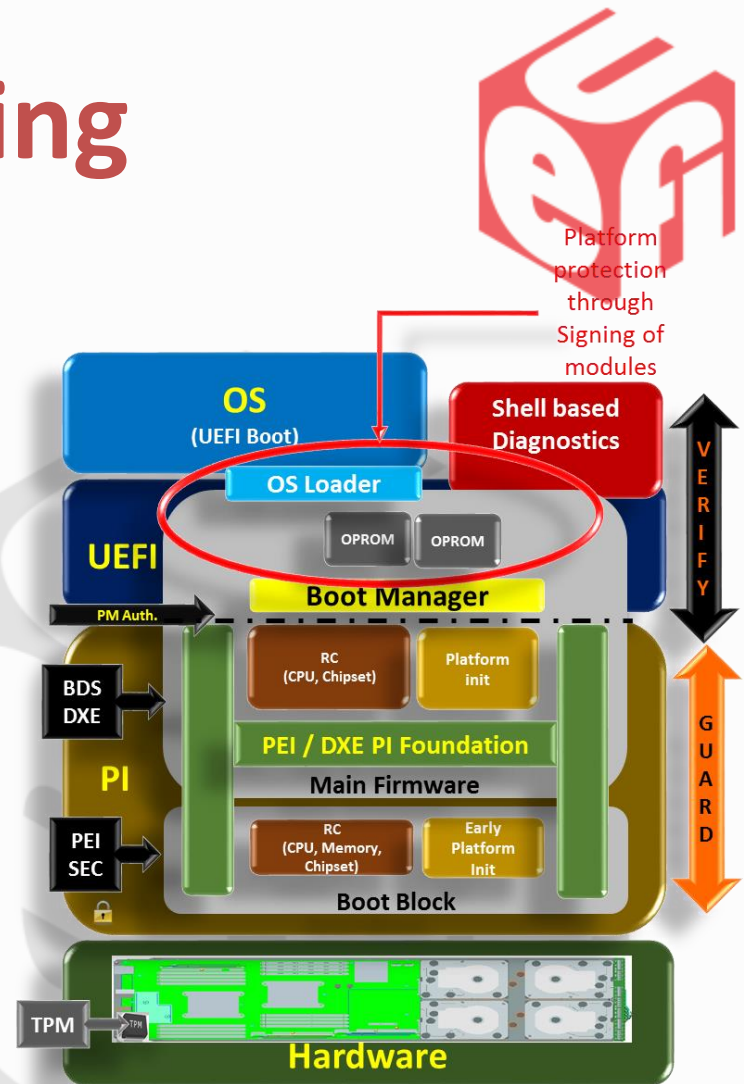


- Signed capsule updates
- UEFI Secure boot
 - local / network
- TPM on the platform
 - Measured boot
 - Root of Trust for Reporting
 - Storage
- Protect machine configuration & UEFI Secure boot trust anchors
- In-band and out-of-band network security

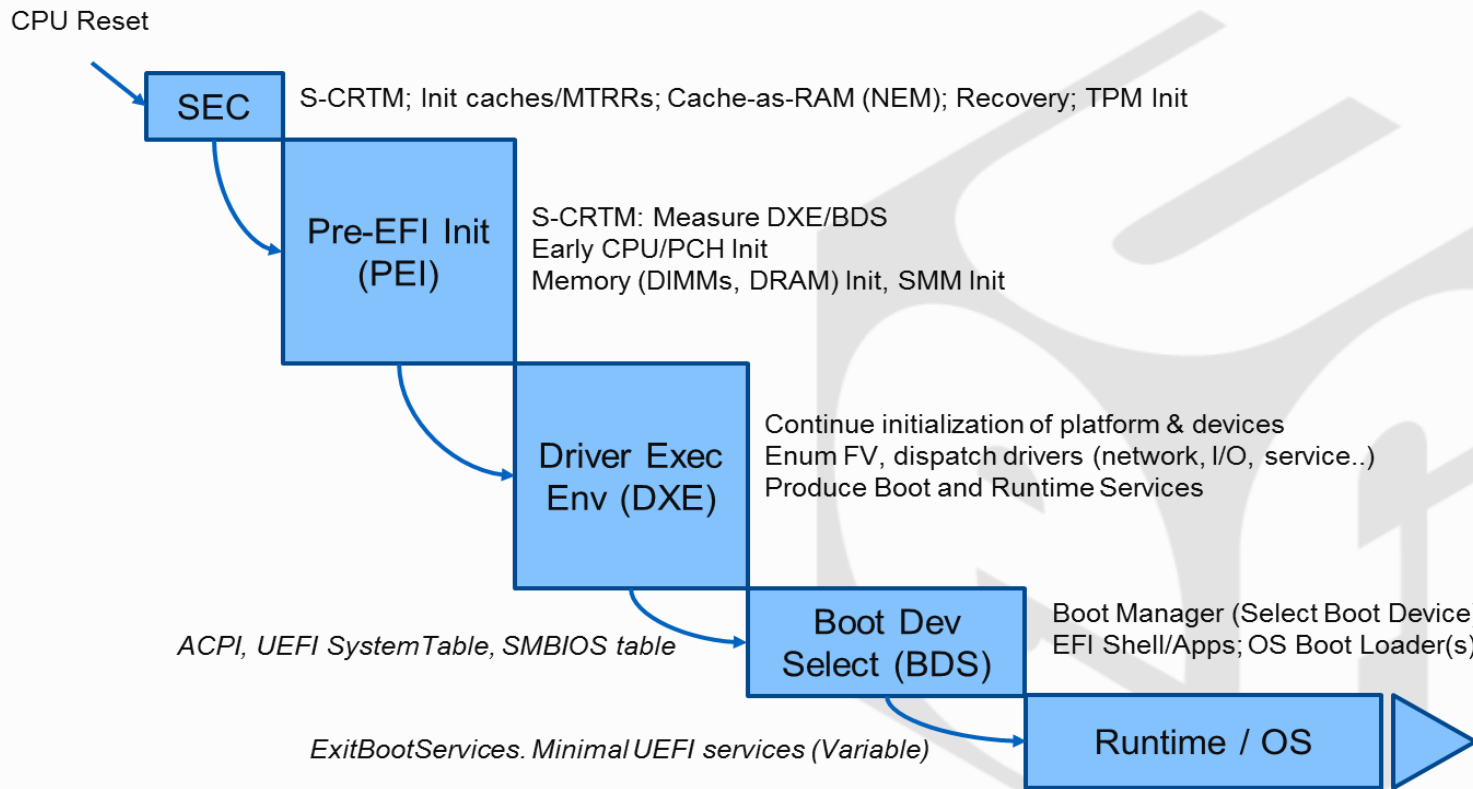


Guarding and Verifying

- PI & UEFI complement each other to impart **platform security** through guarding and verification during pre-boot.
- PI facilitates **platform hardening** by guarding internal firmware ingredients that consume reset vector, initialization of CPU, Memory, Chipset etc.
- UEFI signing allows **robust platform scaling** through verified inclusion of external firmware ingredients such as OPROMS into the trust chain



Full UEFI Boot Flow





Tools and Diagnostics



Tools and diagnostics challenges

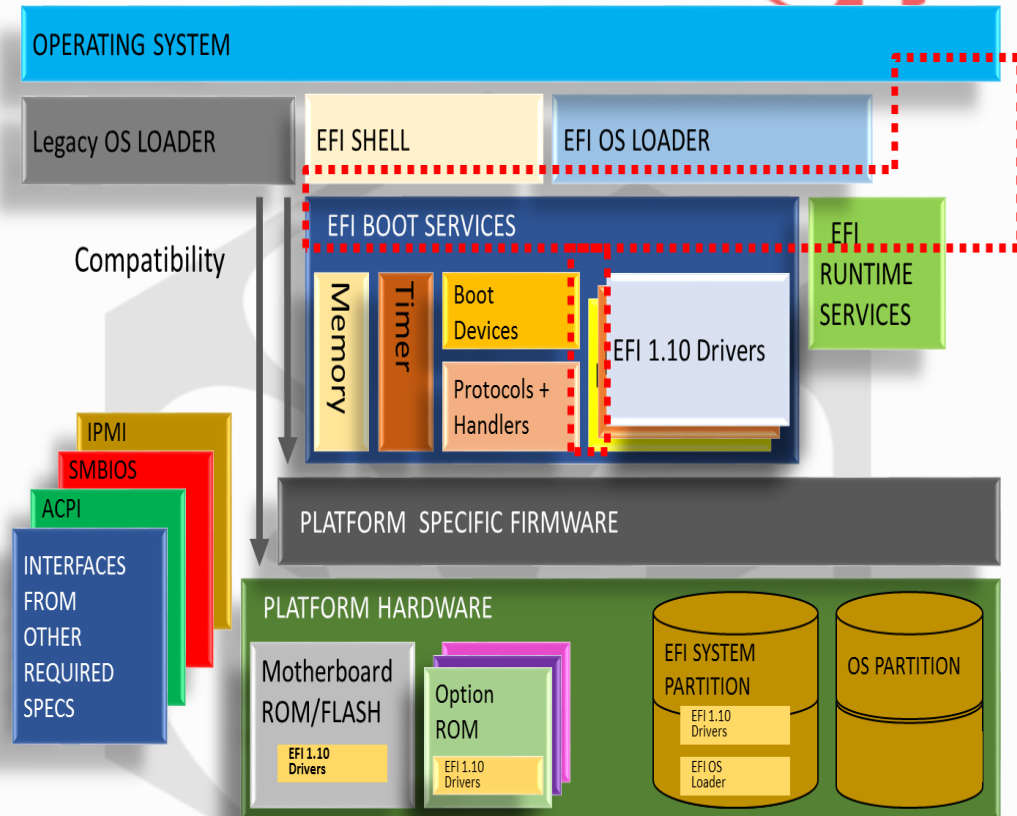
- Platform ingredients from many vendors
- How to assess health, security, compliance of the elements
- Consistent environment to run diagnostics
 - Log / Report / Journal results
- Recovery agent considerations
 - Local / Remote / In-band / Out-of-band



Tools solutions



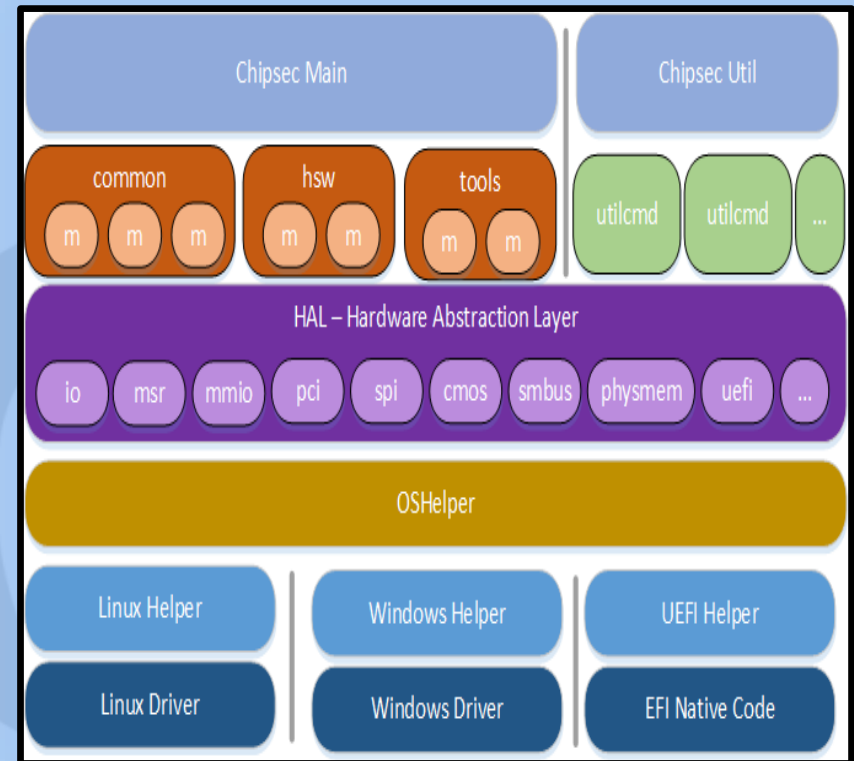
- Environment for hosting tools
 - UEFI Shell
 - Linux UEFI Validation project
- Tools for deployment
 - UEFI SCT
 - PI SCT
 - ACPI Compliance
 - SMBIOS Compliance
 - Security
 - Chipsec
 - Copernicus
 - Selftest



Chipsec tool



- Platform security assessment framework for risk assessment
- Can be extended to meet specific platform security concerns
- Open sourced
 - <https://github.com/chipsec/chipsec>
- Supported Environments
 - Windows
 - Linux
 - UEFI (over Python)



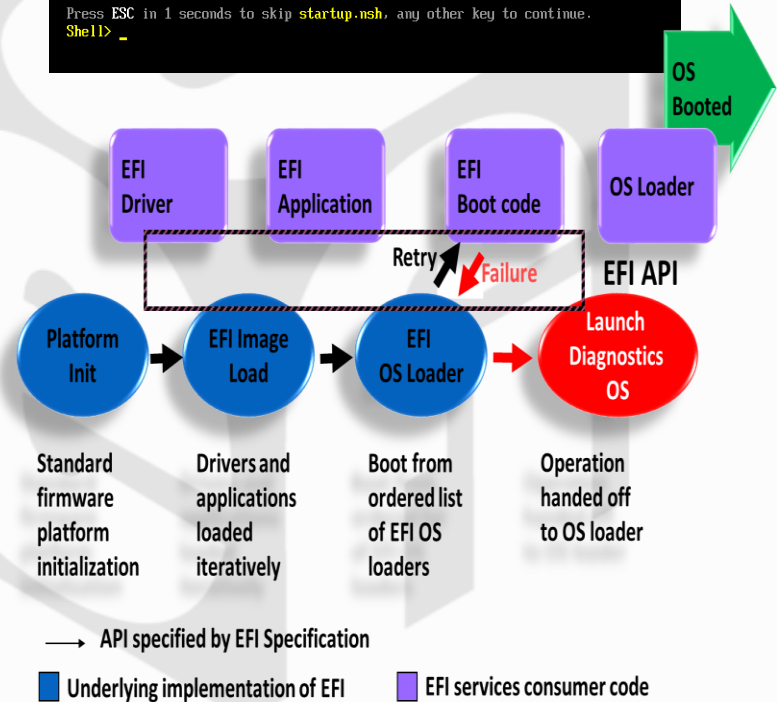
Diagnostic solutions

- Once in UEFI, how to assess, probe, and prod the system
 - Type15 SMBIOS Records
 - Dmpstore for UEFI variables, incl WHEA variable
 - ACPI CA for executing/dumping/viewing namespace
 - UEFI shell to run above, redirect output to file or 'virtual file' (e.g., volatile variable)
 - PCI command to read/write/assess hardware state. Scriptable too
 - Results can be installed in UEFI system table like other hand-off info, or variable, or file on ESP, or sent across the network using UEFI network stack

```

EFI Shell version 2.00 (4096.11)
Current running mode 1.1.2
Device mapping table
fs0      :Removable HardDisk - Alias hd52g0b blk0
         Acpi (PNP0A03.0) /Pci (1D17) /Usb (6,0) /HD (Part1.Sig90909090)
blk0     :Removable HardDisk - Alias hd52g0b fs0
         Acpi (PNP0A03.0) /Pci (1D17) /Usb (6,0) /HD (Part1.Sig90909090)
blk1     :HardDisk - Alias (null)
         Acpi (PNP0A03.0) /Pci (1F12) /Ata (Primary,Master) /HD (Part1.SigD5BAE38B)
blk2     :HardDisk - Alias (null)
         Acpi (PNP0A03.0) /Pci (1F12) /Ata (Primary,Master) /HD (Part2.SigD5BAE38B)
blk3     :BlockDevice - Alias (null)
         Acpi (PNP0A03.0) /Pci (1F12) /Ata (Primary,Master)
blk4     :BlockDevice - Alias (null)
         Acpi (PNP0A03.0) /Pci (1F12) /Ata (Secondary,Master)
blk5     :Removable BlockDevice - Alias (null)
         Acpi (PNP0A03.0) /Pci (1D17) /Usb (6,0)

Press ESC in 1 seconds to skip startup.nsh, any other key to continue.
Shell> _
    
```



Can we do more?



- Yes
- Working group in OCP on updates/management
- Liaison
- Group subteam in UEFI Forum?
- More open source oppty and collaboration

Call To Action



- Get involved in the cloud
- Talk to Mallik and Vincent about how to do more in OCP and the UEFI Forum for Cloud
- The best ideas come from the people who do the work everyday.

More information



- www.opencompute.org – OCP specs
- www.uefi.org – UEFI, ACPI, Shell, PI Specifications
- www.tianocore.org – open source UEFI
- <http://firmware.intel.com> – white papers, training
- chipsec
<https://github.com/chipsec/chipsec>


Thanks for attending the
UEFI Spring Plugfest 2015



For more information on
the Unified EFI Forum and
UEFI Specifications, visit
<http://www.uefi.org>



presented by



Microsoft

