

Impact of platform firmware on Linux kernel

Megha Dey, Sai Praneeth Prakhya
Intel Open Source Technology Center


AGENDA



- Introduction to UEFI
- UEFI implementation bug crashing Linux
- Linux EFI subsystem bug affecting system stability
- Introduction to Linux UEFI Validation (LUV) project
- Make your UEFI system Linux ready with LUV

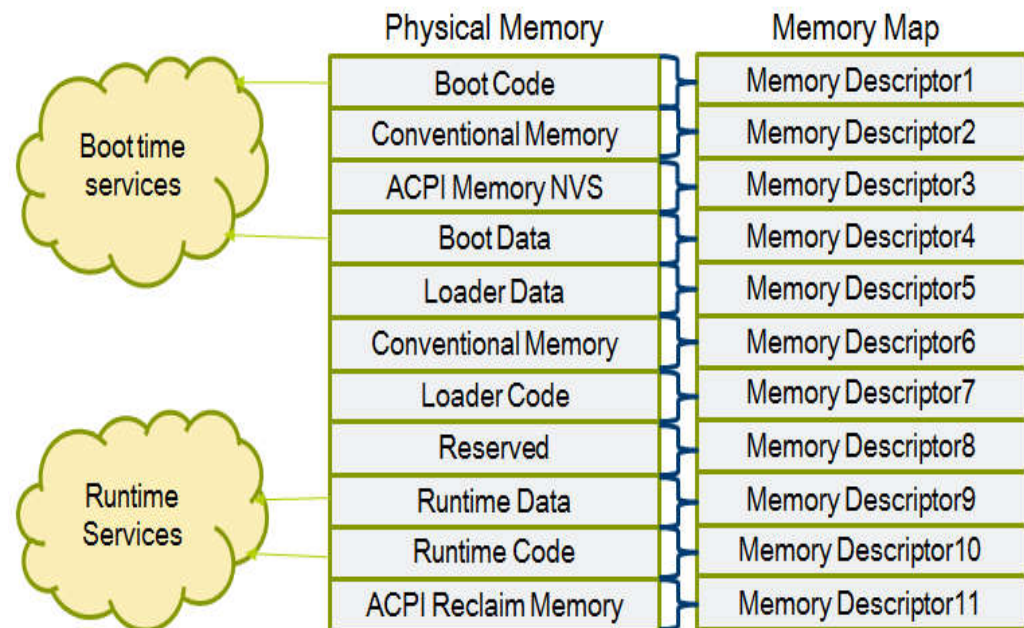
Introduction to UEFI

- UEFI forum members – Many biggies like AMD, Apple, Dell, IBM, Microsoft etc.. and of course Intel.
- UEFI firmware, not only in PC and servers but also in embedded space.
- UEFI (Unified Extensible Firmware Interface) is a specification that defines a software interface between an OS and platform firmware.



UEFI Services

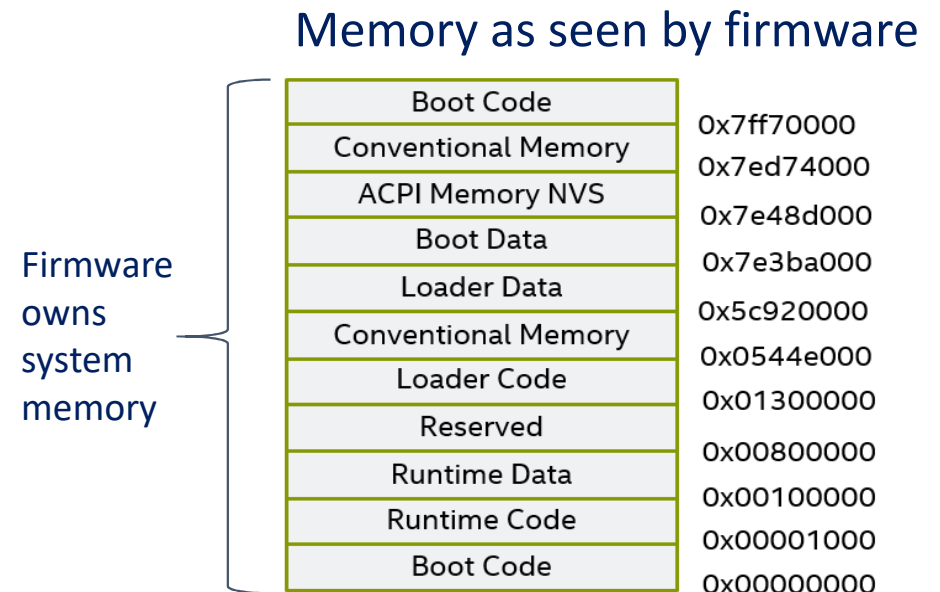
- Implemented as API's – has code and data regions in memory.
- UEFI offers two types of services:
 1. Boot time services: API's available to OS loader before ExitBootServices()
 2. Runtime services: API's available to OS loader **before and after** ExitBootServices()




Boot Process – Before ExitBootServices



- Firmware controls system memory
- Memory accesses in physical mode



Boot Process – After ExitBootServices



- Kernel controls system memory
- Physical addressing mode
- Before calling SetVirtualAddressMap()

Memory as seen by kernel


Free	Boot Code	0x7ff70000
	Conventional Memory	0x7ed74000
	ACPI Memory NVS	0x7e48d000
	Boot Data	0x7e3ba000
	Loader Data	0x5c920000
	Conventional Memory	0x0544e000
Reserved	Loader Code	0x01300000
	Reserved	0x00800000
	Runtime Data	0x00100000
Free	Runtime Code	0x00001000
	Boot Code	0x00000000

Boot Process – Virtual Address space

Kernel Virtual Address space

End Modules	0xffffffff000000
Modules	0xffffffffa0000000
High Kernel Mapping	0xffffffff80000000
EFI Runtime Services	0xffffffffef00000000
ESPfix Area	0xffffffff0000000000
vmemmap	0xffff818000000000
vmalloc() area	0xffff810000000000
Low Kernel mapping	0xffff808000000000
Kernel space	0xffff800000000000
User space	0x0000000000000000

Unused
EFI Runtime Code
EFI Runtime Data
Unused
EFI Runtime Data
EFI Runtime Code


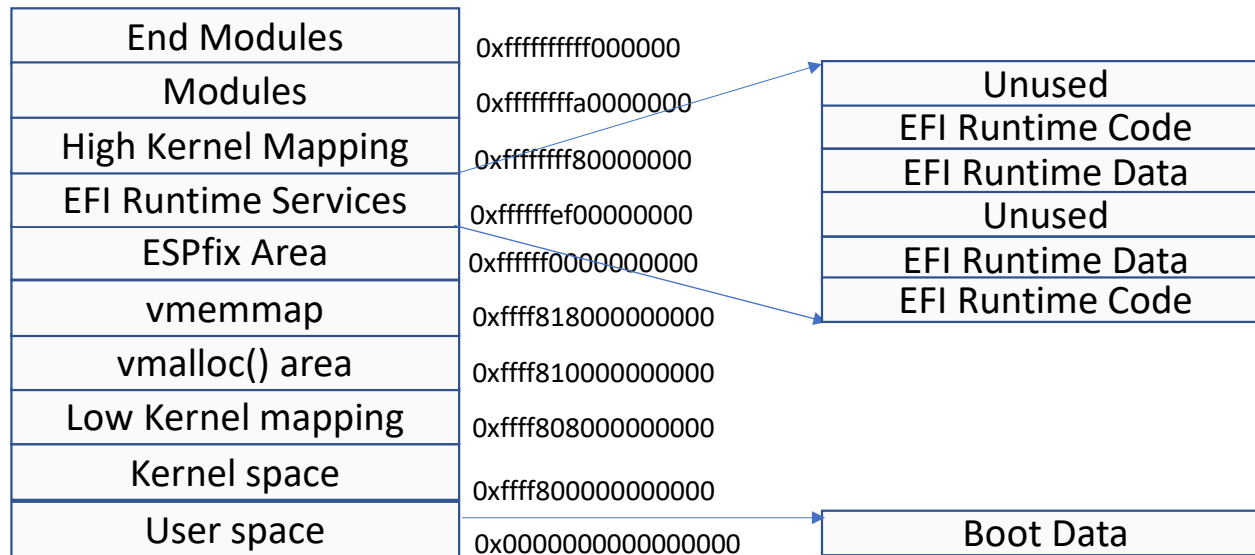


Memory as seen by firmware

Boot Code	0x7ff70000
Conventional Memory	0x7ed74000
ACPI Memory NVS	0x7e48d000
Boot Data	0x7e3ba000
Loader Data	0x5c920000
Conventional Memory	0x0544e000
Loader Code	0x01300000
Reserved	0x00800000
Runtime Data	0x00100000
Runtime Code	0x00001000
Boot Code	0x00000000

Illegal accesses by firmware

Kernel Virtual Address space



Memory as seen by firmware

Boot Code	0x7ff70000
Conventional Memory	0x7ed74000
ACPI Memory NVS	0x7e48d000
Boot Data	0x7e3ba000
Loader Data	0x5c920000
Conventional Memory	0x0544e000
Loader Code	0x01300000
Reserved	0x00800000
Runtime Data	0x00100000
Runtime Code	0x00001000
Boot Code	0x00000000

- EFI Runtime Service that caused illegal access

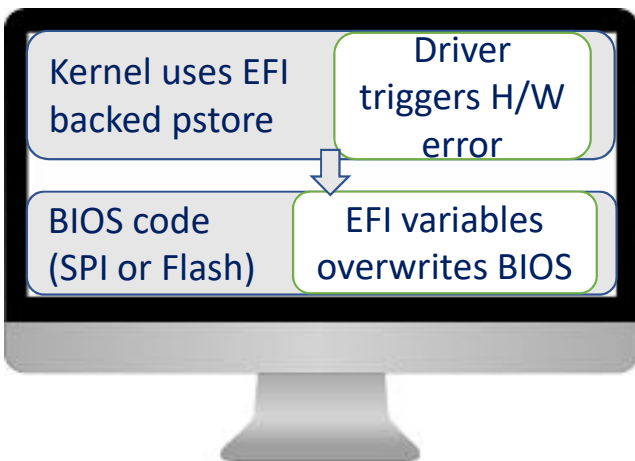
➤ Illegally accessed address

Unmapped region

EFI Runtime Service
referencing EFI
Conventional
Memory

Samsung laptops bricked by UEFI

PC enthusiast reports “Booting Linux bricks Samsung laptops”



Kernel hacker starts investigation



Buggy firmware



How Much is too much

Kernel developer

Reference: <https://mjg59.dreamwidth.org/22855.html>

Kernel bugs causing panic

Kernel panics on BGRT enabled platform

```
[ 0.251599] Last level dTLB entries: 4KB 64, 2MB 0, 4MB 0, 1GB 4
[ 0.259126] Freeing SMP alternatives memory: 32K (ffffffffff8230e000 - ffffffff8231
[ 0.271803] BUG: unable to handle kernel paging request at ffffffffefce35002
[ 0.279740] IP: [] efi_bgrt_init+0x144/0x1fd
[ 0.286383] PGD 300f067 PUD 0
[ 0.289879] Oops: 0000 [#1] SMP
[ 0.293566] Modules linked in:
[ 0.417252] DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400
[ 0.425350] Stack:
[ 0.427638] ffffffff821bbce0 ffffffff82256900 ffff88016dbaea40 ffffffff82003f40
[ 0.436086] ffffffff821bbce0 ffffffff82003f88 ffffffff8219c0c2 0000000000000000
[ 0.444533] ffffffff8219ba4a ffffffff822622c0 0000000000083000 00000000ffffffff
[ 0.452978] Call Trace:
[ 0.455763] [] efi_late_init+0x9/0xb
[ 0.461697] [] start_kernel+0x463/0x47f
[ 0.467928] [] ? set_init_arg+0x55/0x55
[ 0.474159] [] ? early_idt_handler_array+0x120/0x120
[ 0.481669] [] x86_64_start_reservations+0x2a/0x2c
[ 0.488982] [] x86_64_start_kernel+0x13d/0x14c
[ 0.495897] Code: 00 41 b4 01 48 8b 78 28 e8 09 36 01 00 48 85 c0 48 89 c3 75 13
[ 0.518151] RIP [] efi_bgrt_init+0x144/0x1fd
[ 0.524888] RSP []
[ 0.528851] CR2: ffffffffefce35002
[ 0.532615] ---[ end trace 7b06521e6ebf2aea ]---
[ 0.537852] Kernel panic - not syncing: Attempted to kill the idle task!
```



Kernel developer



LUV team

Enable all possible
tables/features in OVMF and use
Qemu + OVMF to find bugs

Reference: <https://lkml.org/lkml/2015/12/10/599>
BGRT - Boot Graphics Resource Table
OVMF - OpenSource Virtual Machine Firmware

Buggy firmware causes issues

ESRT parsed successfully

```
0.000000] efi: EFI v2.60 by EMK II
0.000000] efi: SMBIOS=0x7f1000 ACPI=0x7f1000 ACPI 2.0=0x7f1000
ESRT=0x7f1000 MEMATTR=0x7f1000
0.000000] SMBIOS 2.8 present.
0.000000] DMI: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 0.0.0 02/06/2015
0.000000] e820: last_pfn = 0x7ff70 max_arch_pfn = 0x400000000
0.000000] x86/PAT: PAT not supported by CPU.
0.000000] e820: Reserving ESRT space from 0x000000007f1000 to 0x000000007f1000
0.000000] Scanning 1 areas for low memory corruption
```

ESRT supporting platform firmware

Same kernel booted on two different machines

Failed to parse ESRT

```
[ 0.000000] NX (Execute Disable) protection: active
[ 0.000000] efi: EFI v2.50 by EMK II
[ 0.000000] efi: SMBIOS=0x7f1000 SMBIOS 3.0=0x7f1000 ACPI=0x7f1000
ACPI 2.0=0x7f1000 ESRT=0x7f1000 FRCP=0x7f1000
[ 0.000000] SMBIOS 3.0.0 present.
[ 0.000000] e820: last_pfn = 0x7ff70 max_arch_pfn = 0x400000000
[ 0.000000] efi: requested map not found.
[ 0.000000] esrt: ESRT header is not in the memory map.
[ 0.000000] Scanning 1 areas for low memory corruption
```

ESRT supporting platform firmware

It's tough to catch these errors



QA Engineer

No worries! We have a solution



LUV team

Enable all possible tables/features in OVMF and use Qemu + OVMF to find bugs


ESRT - EFI System Resource Table

UEFI implementation is critical

Things can go wrong during interactions between firmware and Linux kernel

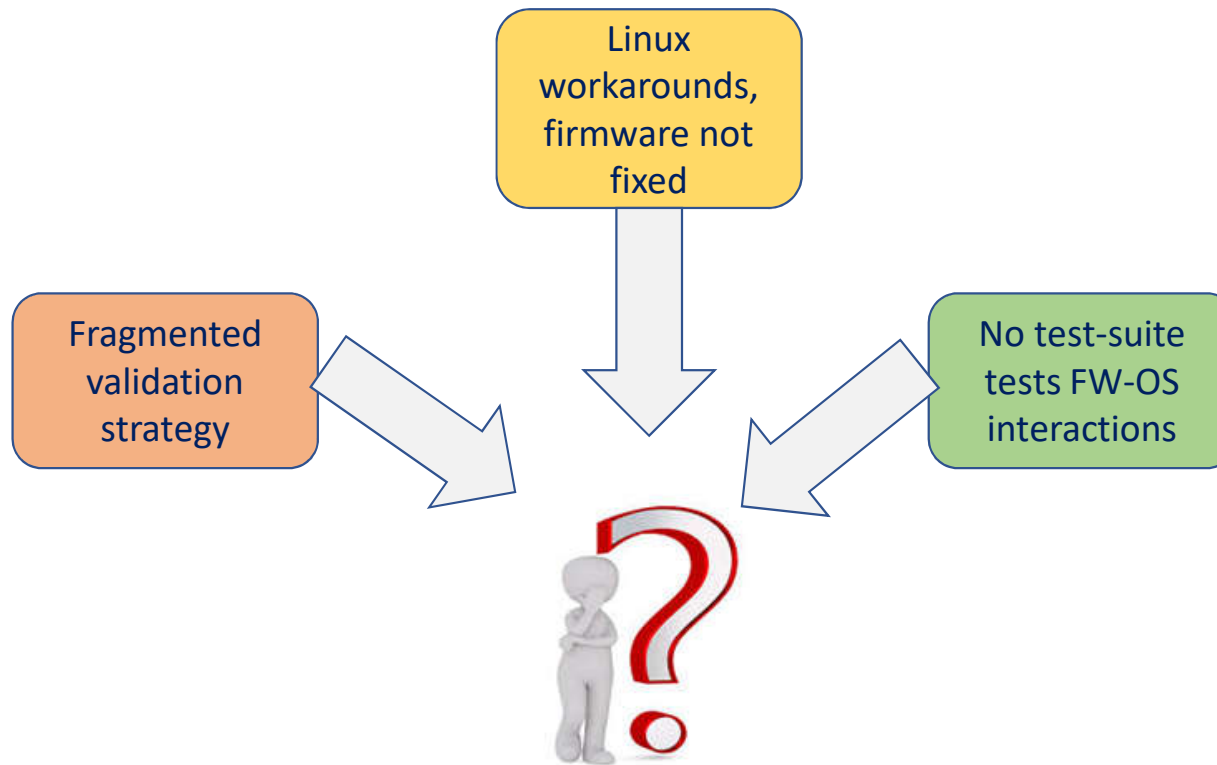


Faulty implementation of generic UEFI hurts Linux

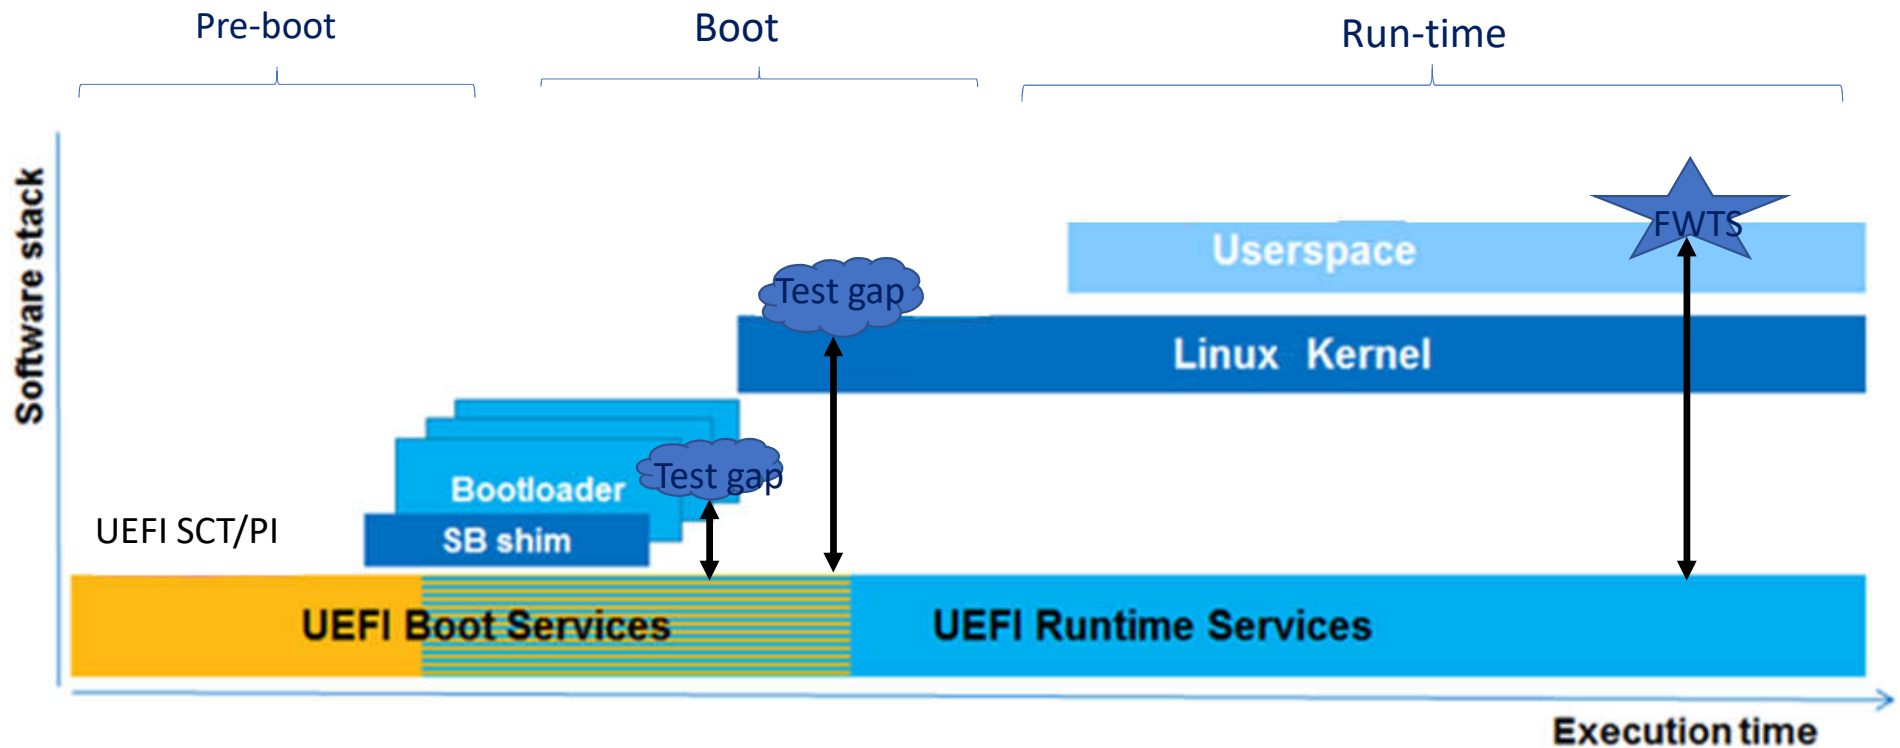


Linux UEFI Validation (LUV) project was started

Why LUV?



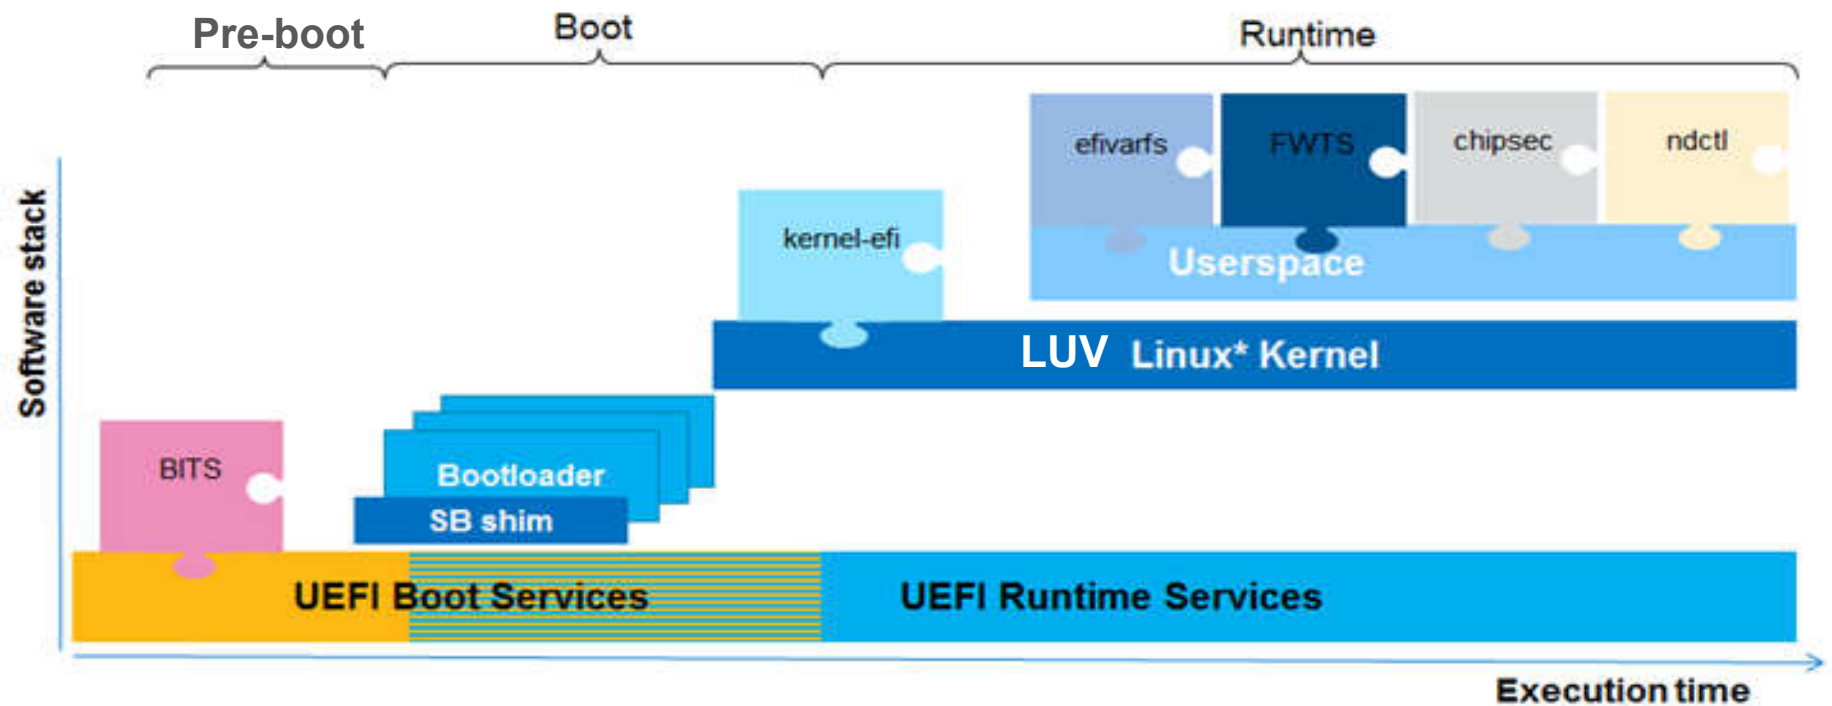
Current Validation Strategy





-
- efivarfs
- BITS
- chipsec
- ndctl
- FWTS
- Kernel-efi
- Your test

LUV: Covering the entire spectrum



LUV: Detect bugs early

At boot time:

```
Linux UEFI Validation Distribution 2.1-rc2
Date and time of the system : 2017-03-28--13-15-48
System: Intel(R) Core(TM) i7-7820HQ CPU @ 2.90GHz x 8

Test results summary:
[+]bits
[+]chipsec
[+]efivarfs
[+]fwts
[-]kernel-efi-warnings
  [-]EFI_illegal_accesses_test... 1 failures!
    Test: EFI_illegal_accesses_test
    Description: Check for illegal accesses to EFI memory regions.
    FAIL: Found 1 instances of "efi: \[Firmware Bug\]: Fixing illegal access to EFI region at PA: ".
    Related kernel information
    -----
    [ 0.012085] pid_max: default: 32768 minimum: 301
    [ 0.017353] ACPI: Core revision 20160930
    [ 0.060109] ACPI: 17 ACPI AML tables successfully acquired and loaded
    [ 0.067511] efi: [Firmware Bug]: Fixing illegal access to EFI region at PA: 0xb82ffc98
    [ 0.076566] --- [ User Space ] ---
    [ 0.080243] 0x0000000000000000-0x00000000000001000 4K RW GLB x pte
    [ 0.090241] 0x00000000000001000-0x00000000000200000 2044K pte
    -----
  [+][Kernel_FW_BUG... 1 failures!
  [+][Kernel_FW_WARN... passed
  [+][Kernel_FW_INFO... passed
[+]indctl

Ran 6 testsuites and 134 unittests, 72 passes, 268 fails, 44 skipped, 6 warnings.
```

Kernel fixing illegal access
by firmware during
Set Virtual Address Map

At run-time:

```
[ 238.856558] [-] fwts
[ 238.859017] [+] version... skipped
[ 238.863030] [+] bios_info... passed
[ 238.867136] [+] prd_info... skipped
[ 238.871245] [+] oops... passed
[ 238.874859] [+] olog... skipped
[ 238.974803] [+] klog... 1 failures!
[ 238.978925] [+] kmc_info... skipped
[ 239.031610] [+] mtrr... 3 failures!
[ 239.036934] [+] acpiinfo... passed
[ 239.044758] efi: [Firmware Bug]: Fixing illegal access to EFI region at PA: 0xacc2520
[ 239.053614] efi: [Firmware Bug]: Fixing illegal access to EFI region at PA: 0xad7a5cd6
```

Kernel fixing illegal access by firmware
when EFI Runtime Service is invoked


LUV uses hardened kernel to
work around these issues
(during boot time and run
time)

LUV: Easy to use

Plug and play



Diskless boot

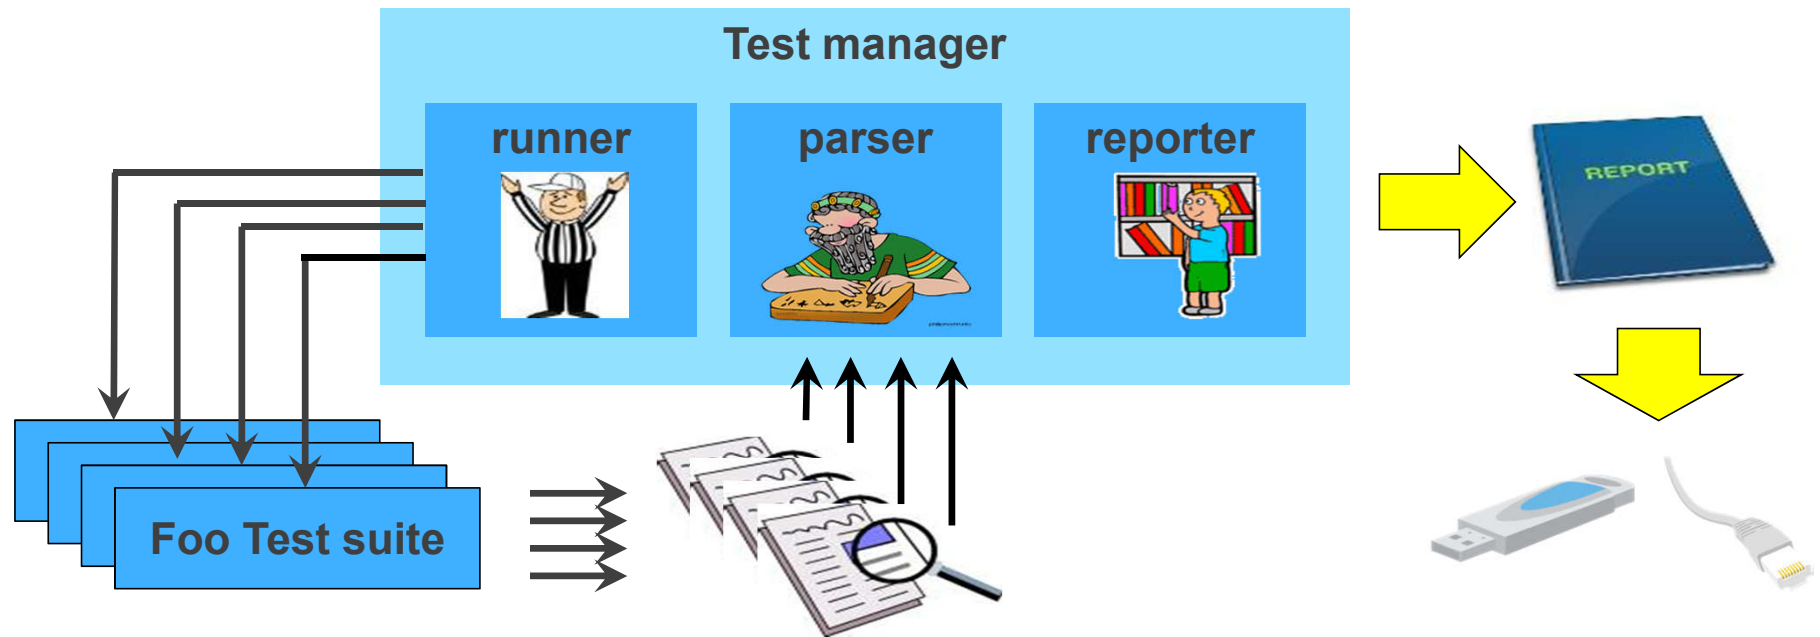


Boot on Virtual Machine



LUV: Easy to contribute

LUV is open source!




LUV: Available features

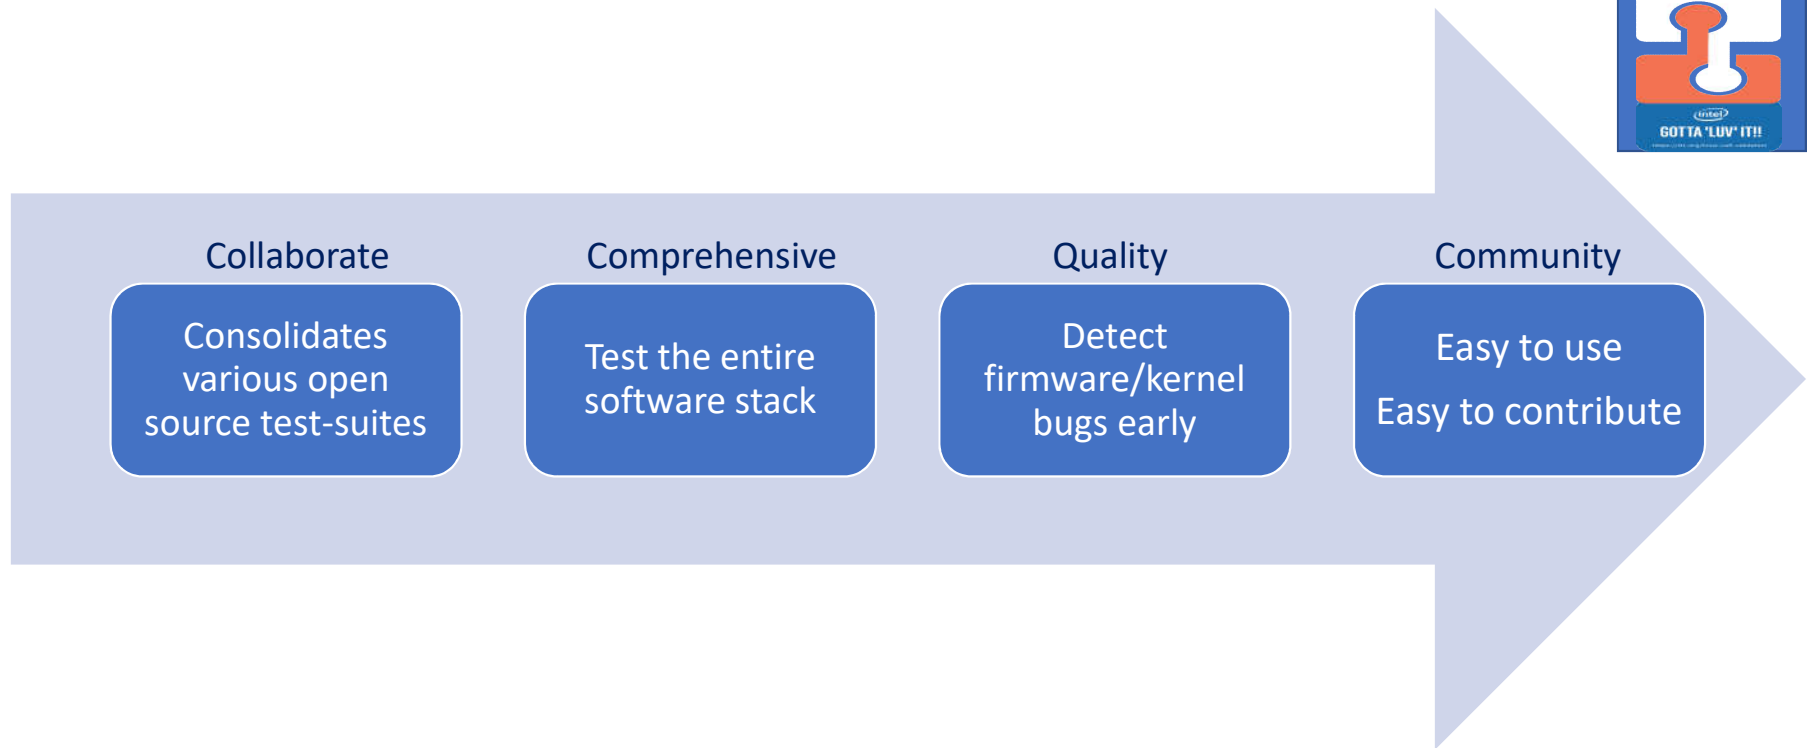
- Results are saved for viewing
 - If (USB Stick) - Saved in LUV_RESULTS folder
 - If (Net Boot) - Saved to HTTP server
 - If (VM) - Saved in the results partition of luv.img
- Debug via network using Netconsole.
- Telemetry mechanism to send error report to the LUV server incase of crash.
- Lists all available DSMs on a platform.
- *Future feature: Add support for more bootloaders.

```
Linux UEFI Validation Distribution 2.2-dev Date and time of the system : 2018-02-15--21-39-35
System: Intel(R) Core(TM) i7-4770HQ CPU @ 2.20GHz x 8 Test results summary:
chipsec
+chipsec.modules.common.bios_ts... passed
+chipsec.modules.common.spi_access... 2 warnings
+chipsec.modules.common.ia32cfg... passed
+chipsec.modules.common.spi_lock... passed
+chipsec.modules.common.smm... passed
+chipsec.modules.common.bios_kbrd_buffer... passed
+chipsec.modules.common.bios_smi... passed
+chipsec.modules.common.smmr... passed
+chipsec.modules.common.spi_fdopss... passed
+chipsec.modules.common.spi_desc... passed
+chipsec.modules.common.rtclock... 1 warnings
+chipsec.modules.common.bios_wp... 1 failures!
+chipsec.modules.common.secureboot.variables... 1 warnings
+chipsec.modules.common.uefi.s3bootscript... 1 failures!
+chipsec.modules.common.uefi.access_uefispec... passed
+chipsec.modules.memconfig... passed
+chipsec.modules.remap... passed
+chipsec.modules.smm_dma... passed
efivarfs
fwts
kernel-efi-warnings
indctl
pstore-tests

Ran 6 testsuites and 159 unittests, 80 passes, 117 fails, 59 skipped, 4 warnings.
LEGEND:
Severity of failures: Low: [red circle] Medium: [red circle] High: [red circle] Unknown: [red circle]
```



LUV: Path to better firmware



What do you need?

Tell us your needs to better validate your firmware!



Join US

<https://01.org/linux-uefi-validation>

Credits


Neri Ricardo

Matt Fleming

Gayatri Kammela

Naresh Bhat

And many more contributors...





Thank you.

Any Questions?