BARing the System New vulnerabilities in Coreboot & UEFI based systems

Presenting: Yuriy Bulygin (@c7zero), Oleksandr Bazhaniuk (@ABazhaniuk) Andrew Furtak, John Loucaides, Mikhail Gorobets



Agenda

- Recap of SMM Pointer Vulnerabilities
- Intro to Memory-Mapped I/O
- MMIO BAR Issues
- MMIO BAR Issues in UEFI Firmware
- MMIO BAR Issues in Coreboot Firmware
- Limitations
- Mitigations
- Tools
- Conclusion

Recap of SMM pointer vulnerabilities

Pointer Arguments to SMI Handlers



Exploiting SMM pointers...



Exploit tricks SMI handler to write to an address in SMRAM (Attacking and Defending BIOS in 2015)

Attacking hypervisors via SMM pointers...



Even though SMI handler check pointers for overlap with SMRAM, exploit can trick it to write to VMM protected page (Attacking Hypervisors via Firmware and Hardware)

Example: SMIFlash SMI Handler

VOID SMIFlashSMIHandler (2 . . 3 SwSmi = (UINT8)DispatchContext->SwSmiInputValue; CpuState = pSmst->CpuSaveState; 4 AddrHi = CpuState[Cpu].Ia32SaveState.ECX; 5 AddrLo = CpuState[Cpu].Ia32SaveState.EBX; 6 7 Buff = AddrHi; Buff = Shl64(Buff, 32);8 Buff += AddrLo; 9 10 . . . 11 switch (SwSmi) { 12 . . 13 case 0x21: ReadFlashData((FUNC BLOCK *) Buff); 14 15 . . 16 case 0x25: ReadFlashInfo((INFO BLOCK *)Buff); 17 18 19 EFI STATUS ReadFlashData(IN OUT FUNC BLOCK *func) 20 21 . . 22 sts = Flash->Read(23 (UINT8*) (FlashStart + func->BlockAddr), func->BlockSize, 24 (UINT8*) func->BufAddr 25 26);

Reported by ATR to BIOS vendor in June 2014

Similar to publication by Sogeti ESEC Lab

void EFIAPI SwSMIDispatchFunction(EFI_HANDLE DispatchHandle, EFI_SMM_SW_DISPATCH_CONTEXT

```
// ...
```

```
struct smiflash_arg *pointer; // see bellow for the struct
int smi number = DispatchContext->SwSmiInputValue;
```

```
// Retrieve a pointer from user provided value
pointer = ecx << 32 | ebx;</pre>
```

```
if (smi_number != 0x25)
    pointer->unknown = 0;
```

```
// some init ...
```

```
switch (smi_number) {
    case 0x20:
        // ...
        break;
    case 0x21:
        swsmi handler21(pointer);
```

SMI handlers now validate input pointers

- SMI handlers now validate pointer + offsets received from the OS for overlap with SMRAM before using it (SmmIsBufferOutsideSmmValid). This does not block exploits using SMI handlers as proxies to attack hypervisor pages (Hyper-V, Windows 10 Virtual Secure Mode)
- Most recently, EDKII implemented CommBuffer at fixed memory location to mitigate attacks on hypervisors and reporting to Windows through the Windows SMM Mitigations ACPI Table (WSMT)

Length	Bit offset	Description
1	0	FIXED_COMM_BUFFERS If set, expresses that for all synchronous SMM entries, SMM will validate that input and output buffers lie entirely within the expected fixed memory regions.
1	1	COMM_BUFFER_NESTED_PTR_PROTECTION If set, expresses that for all synchronous SMM entries, SMM will validate that input and output pointers embedded within the fixed communication buffer only refer to address ranges that lie entirely within the expected fixed memory regions.
1	2	SYSTEM_RESOURCE_PROTECTION If set, expresses that firmware has taken steps ensuring that configuration for any system resources that are not configurable through an architectural mechanism (e.g. ACPI or PCI MMIO) must be locked by firmware before transitioning to Windows.

Memory-Mapped I/O (MMIO)

PCI Express

- PCI Express Fabric consists of PCIe components connected over PCIe interconnect in a certain topology (e.g. hierarchy)
- Root Complex is a root component in a hierarchical PCIe topology with one or more PCIe root ports
- Components: *Endpoints* (I/O Devices), *Switches*, PCIe-to-PCI/PCI-X *Bridges*
- All components are interconnect via PCI Express Links
- Physical components can have up to 8 physical or virtual functions
- Some endpoints are *integrated* into Root Complex

PCIe Config Space Layout



OM14301A

Figure 7-3: PCI Express Configuration Space Layout

Source: PCI Express Base Specification Revision 3.0

PCI/PCIe Config Space Access

1. Software uses processor I/O ports CF8h (control) and CFCh (data) to access PCI configuration of bus/dev/fun. Address (written to control port) is calculated as:



- 2. Enhanced Configuration Access Mechanism (ECAM) allows accessing PCIe extended configuration space (4kB) beyond PCI config space (256 bytes)
 - Implemented as memory-mapped range in physical address space split into 4kB chunks per B:D.F
 - Register address is a memory address within this range

MMCFG base + bus*32*8*1000h + dev*8*1000h + fun*1000h + offset

Memory-Mapped I/O

- Devices need more space for registers
- → Memory-mapped I/O (MMIO)
- MMIO range is defined by Base Address Registers (BAR) in PCI configuration header
- Access to MMIO ranges forwarded to devices





MMIO BARs

MMIO Range		BAR		Base		Size		En		Description
GTTMMADR SPIBAR HDABAR GMADR DMIBAR MMCFG RCBA	 	00:02.0 + 10 00:1F.0 + F0 00:03.0 + 10 00:02.0 + 18 00:00.0 + 68 00:00.0 + 60 00:1F 0 + F0	 	00000000000000000 0000000000000000 00000	 	0040000 0000200 00001000 00001000 00001000 00001000	 	 1 1 1 1 1 1 1	 	Graphics Translation Table Range SPI Controller Register Range HD Audio Register Range Graphics Aperture Root Complex Register Range PCI Express Register Range PCH Root Complex Register Range
MCHBAR		00:00.0 + 48		000000000FED10000		00008000		1		Memory Controller Register Range

•••

MMIO Range Relocation

- MMIO ranges can be *relocated* at runtime by the OS
 - OS would write new address in BAR registers

- Certain MMIO ranges cannot be relocated at runtime
 - Fixed (e.g. direct-access BIOS range)
 - Or locked down by the firmware (e.g. MCHBAR)

MMIO BAR Issues

Firmware use of MMIO



MMIO BAR Issue

Exploit with PCI access can modify BAR register and relocate MMIO range

On SMI interrupt, SMI handler firmware attempts to communicate with device(s)

It may read or write "registers" within relocated MMIO



Examples of MMIO BARs accessed in SMM

- EHCI (USB 2.0) controller MMIO BAR (B0:D26:F0, B0:D29:F0)
- GBe LAN MMIO BAR (B0:D25:F0)
- Root Complex Block Address (RCBA) on earlier platforms (B0:D31:F0)
- SPI BAR on Skylake or later generations (B0:D31:F5)
- AHCI (SATA) controller MMIO BAR (B0:D31:F2, B0:D31:F5)
- xHCI (USB 3.0) controller MMIO BAR (B0:D20:F0)
- Integrated Graphics Device MMIO BAR (B0:D2:F0)
- B1:D0.F0 MMIO BAR

SPI Controller MMIO BAR (Access to SPI Flash)

chipsec_util.py uefi var-write B 5555555-4444-3333-2211-000000000000 B.bin
chipsec util.py mmio dump SPIBAR

[CHIPSEC] Dumping SPIBAR MMIO space.. [mmio] MMIO register range [0x00000000FE010000 +00000000: 07FF0200 +00000004: 0000E000 SPI Status and Control +00000008: 002558AC +0000000C: 00000000 +00000010: 4242423F SPI Flash Address (address +00000014: 42424242 variable is written to in flash) +00000018: 42424242 +0000001C: 42424242 +00000020: 42424242 +00000024: 42424242 +00000028: 42424242 +0000002C: 42424242 +00000030: 42424242 SPI Flash Data +00000034: 42424242 (Variable contents) +00000038: 42424242 -0000003C: 42424242

Finding MMIO BAR issues at runtime

Goal: Find all MMIO registers modified by SMI handler

- 1. Dump MMIO range
- 2. Trigger SMI
- 3. Dump MMIO range and compare all registers

Problem: many registers are modified by devices all the time! Up to 30,000 registers change in Graphics Device MMIO

Finding MMIO BAR issues at runtime

Goal: Find all MMIO registers modified by SMI handler

- 1. Dump MMIO range multiple times
- 2. Find all registers which frequently change without SMM
- 3. Dump MMIO range
- 4. Trigger SMI
- 5. Dump MMIO range and compare all registers
- 6. Find registers which don't normally change
- 7. Repeat this multiple times to confirm suspected registers are actually being modified in SMM
- 8. Copy original contents of MMIO range to memory
- 9. Relocate MMIO range (change its base address) to this memory
- 10. Generate SMI
- 11. Monitor changes in memory at suspected offsets



MMIO BAR Issues in UEFI Firmware

Finding MMIO BAR issues in binaries

- 1. Consider MMIO BAR (MBARA) of GBe LAN device (B0:D25:F0) at offset 0x10
- 2. Legacy PCIe config address is

 $(25 << 11) + 0 \times 10 = 0 \times C810$

0x8000C810 if with Enable bit (31) set

 Memory-mapped ECAM address is ECAM base + offset to 4kB page of B0:D25:F0 + BAR register offset

 $0 \times F8000000 + 0 \times C8010 = 0 \times F80C8010$

4. Look for these constants in the binaries of SMI handlers

GBe LAN MMIO BAR (B0:D25:F0)



GBe LAN MMIO BAR (B0:D25:F0) Access to unchecked **MBARA MMIO** MemorySetResetValues32((unsigned int *)(unsigned int)(MBARA MMIC + 3856), v8, v7); *(_DWORD *)(unsigned int)(MBARA_MML0 + 32) = 71237632; sub 18000156C(0xFA0ui64, 6144); v0 = 70845440: if (*(_DWORD *)(MBARA_MMI0 + 3856) & 8) v0 = 0x4390440;if (*(_DWORD *)(MBARA MMI0 + 3856) & 4) v0 |= 4u; *(DWORD *)(unsigned int)(MBARA MMLO + 32) = v0; *(_DWORD *)(unsigned int)(<mark>MBARA_MMLO</mark> + 32) = 71263232;

```
v9 = (int *)(unsigned int)(MBARA MMID + 32);
```

EHCI MMIO BAR (B0:D29:F0)



Finding MMIO BAR issues in binaries

Identify functions reading PCI config registers via legacy or ECAM access and find ones reading BAR registers using above constants

Legacy P(nfig read	
<pre>pci_dword_read pci_dword_read</pre>	proc n sub mov call mov add jmp endp	ear rsp, 28h edx, ecx cx, 0CF8h outdword cx, 0CFCh rsp, 28h indword	; CODE XREF: sub_10008450+2E [†] p ; sub_100085C4+B2 [†] p

Access to register 0x88 in B0:D31:F0

```
__int64 sub_10000320()
{
    unsigned __int32 v0; // eax@1
    v0 = pcie_read_dword(0i64, 0x1Fu, 0, 0xB8u);
    return pcie_write_dword(0, 0x1Fu, 0, 0xB8u, v0);
}
```

Extended PCIe config access through MMCFG

```
unsigned __int32 __fastcall pcie_read_dword(__int64 a1, unsigned __int8 a2, unsigned __int8 a3, unsigned __int16 a4)
{
    if ( a4 >= 0x100u )
        return *(_DWORD *)((a3 << 12) + (a2 << 15) + ((unsigned __int8)a1 << 20) + (unsigned int)a4 - 0x8000000);
    outdword(0xCF8u, a4 & 0xFC | ((a3 | 8 * (a2 | 32 * (unsigned __int8)a1)) << 8) | 0x80000000);
    return indword(0xCFCu);
}
```

MMIO BAR Issues in Coreboot Firmware

Finding MMIO BAR issues in the source code

- 1. Find functions within SMI handlers which read MMIO BAR PCI config registers (offsets 0x10-0x24 or chipset specific offsets for integrated devices)
 - BAR registers can be read using memory-mapped config reads (offsets in ECAM memory space). In this case, normal memory reads will be used

```
reg_base = (void *)((uintptr_t)pci_read_config32(SA_DEV_IGD,
PCI_BASE_ADDRESS_0) & ~0xf);
```

2. Find all memory accesses to offsets off of BAR addresses within SMI handlers

```
write32(reg_base + PCH_PP_CONTROL, pp_ctrl); <<< memory write of modified
pp_cntrl value
read32(reg_base + PCH_PP_CONTROL);</pre>
```

3. Can often search the names of the BAR registers and MMIO ranges (e.g. SPI_BARO, RCBA/RCRB, PCI_BASE_ADDRESS etc.)

mainboard_io_trap_handler SMI handler



southbridge_smi_sleep SMI handler

```
static void backlight off (void)
reg base pointer points
     void *reg base;
                                                                              relocateable IGD MMIO
     uint32 t pp ctrl;
     uint32 t bl off delay;
     reg base = (void *) ((uintptr t)pci read config32(SA DEV IGD, PCI BASE ADDRESS 0) & ~0xf);
                                                       SMI handler then uses
     /* Check if backlight is enabled */
                                                     reg base to read-modify-
     pp ctrl = read32 (reg base + PCH PP CONTROL);
                                                     write PP CONTROL register
     if (! (pp ctrl & EDP BLC ENABLE))
          return:
                                                               southbridge smi sleep(void)
     /* Enable writes to this register */
     pp ctrl &= ~PANEL UNLOCK MASK;
                                                       /* Figure out SLP TYP */
     pp ctrl |= PANEL UNLOCK REGS;
                                                       reg32 = inl(ACPI BASE ADDRESS + PM1 CNT);
                                                       printk(BIOS SPEW, "SMI#: SLP = 0x%08x\n", reg32);
     /* Turn off backlight */
                                                       slp typ = (reg32 >> 10) & 7;
     pp ctrl &= ~EDP BLC ENABLE;
                                                       switch (slp typ) {
     write32(reg base + PCH PP CONTROL, pp ctrl);
                                                       case SLP TYP S5:
     read32 (reg base + PCH PP CONTROL);
                                                           printk(BIOS DEBUG, "SMI#: Entering S5 (Soft Power off)\n");
              Vulnerable backlight off is
                                                           /* Turn off backlight if needed */
```

backlight off();

invoked when system goes to S5

2

3

4

5

6

7 8

9

10

11

12

13

14 15

16

17

18

19

20 21

22

Limitations

1. Exploit can overwrite specific offsets off of aligned addresses

- MMIO ranges are typically normally (size) aligned
- Most MMIO ranges are 4kB large (Graphics MMIO is 2-4MB)
- Example: 16kB aligned Root Complex Base + 0x38xx (SPI registers)
- PCI architecture allows MMIO ranges as small as 16 bytes
- 2. Exploit may not be able to control values written
 - Firmware SMI handlers typically write specific values to MMIO registers
 - Often do Read-Modify-Write: reg +/- 0x10
 - Certain SMI handler may write attacker-supplied data
 - SetVariable SMI handler write contents of UEFI variable supplied by the OS to SPI_DATAx registers in SPIBAR MMIO range

Limitations

- 1. Many conditions for SMI handler to start communicating with I/O device/controller
 - Device present/enabled, mode/feature supported
 - Is platform in ACPI mode?
 - Other SMM code may also use fake MMIO (and hang)
 - SMI may get triggered on difficult events power button, on S3 resume, etc.
- 2. Often, SMI handlers implement protocol rather than just reading or writing to MMIO registers
 - IF Bit X in Reg1 is set THEN Write to Reg2
 - Poll until certain bits are set/cleared in MMIO register (wait until SPI cycle complete)
 - When SMI handler waits for the device to respond or cycle to complete then it'll hang after MMIO BAR is relocated
- 3. Non PCI-architectural BAR registers are locked down by boot firmware and cannot be relocated (MCHBAR, DMIBAR etc.)

Mitigations

Option 1. SMI handlers can verify MMIO BAR doesn't overlap with SMRAM

Option 2. Firmware can verify that MMIO BAR is not in DRAM (e.g. between TOLUD and 4GB or above TOUUD). This would ensure all BARs used by firmware are within MMIO

Option 3. Firmware can reserve default MMIO range for all BARs. Before accessing MMIO range, SMI handlers can relocate BARs to the default range if they point to somewhere else

SPIBAR Mitigation Example

- On latest platforms, SPI MMIO is a separate 4kB range rather than a part of Root Complex MMIO
- Firmware reserves 0xFE010000 page for SPI MMIO and programs SPI_BAR0 register in SPI controller with this address.
- On any PCH SMI, SMI handler checks SPI_BAR0 and restores it to 0xFE010000 if it's been relocated.



Tools to assist in finding/analyzing these issues

tools.smm.rogue_mmio_bar

Attempts to create fake MMIO ranges in memory, relocate hardware MMIO BARs to the fake memory, then observe changes made by SMI handlers in relocated MMIO ranges

tools.smm.bar

Simply monitors changes made by SMI handlers in MMIO registers of specified MMIO BARs

Conclusion

- The root cause is that firmware assumes hardware is trusted
- Hardware registers like PCI Base Address Registers can be modified by runtime software (some are locked down)
- Firmware shouldn't assume addresses in BAR registers are correct and should treat them as untrusted input
- Boot firmware should also validate contents of BAR registers upon resume from sleep if it restores them from S3 boot script

Thank You!