

# Is OVMF too slow for Serverless Confidential Computing?

Tobin Feldman-Fitzthum

IBM Research

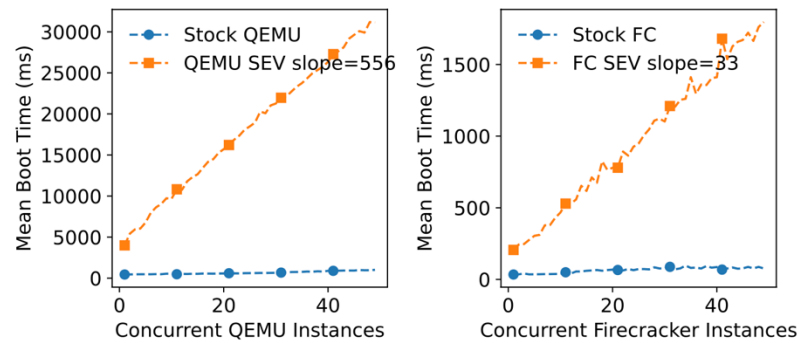
# Disclaimers

- I like OVMF
- I don't like serverless

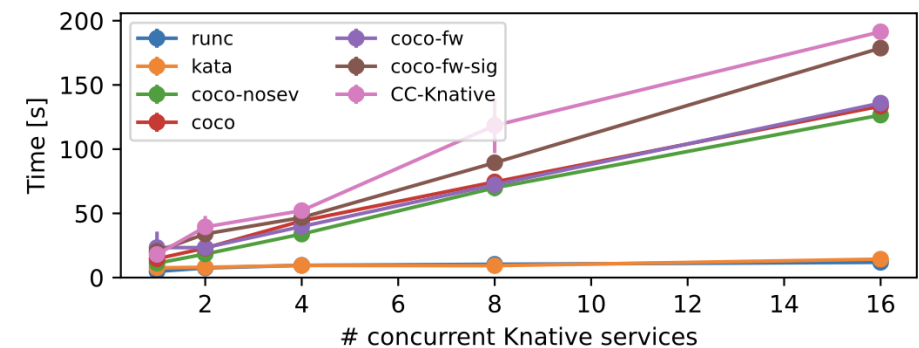


# Serverless Confidential Computing

- Warm starts are slower and more complex
- Cold starts are slower and more complex
- Starting a lot of guests at once is slow




**Figure 12.** Average boot time of concurrent SEV guests from 1 to 50 concurrent instances.



**(b) Throughput-latency of service instantiation.**

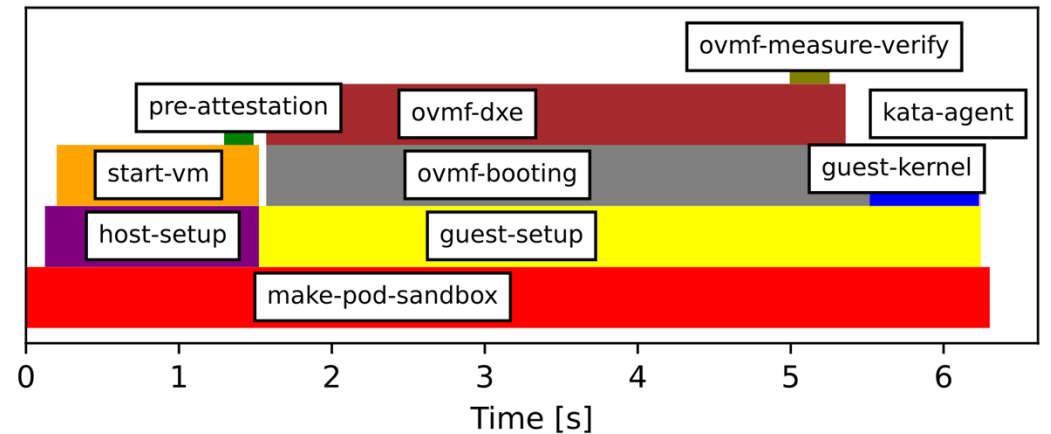
# Cold Starts



**Figure 3.** Breaking down the OVMF boot process with SEV-SNP shows that the boot verifier is a small portion of overall boot time.

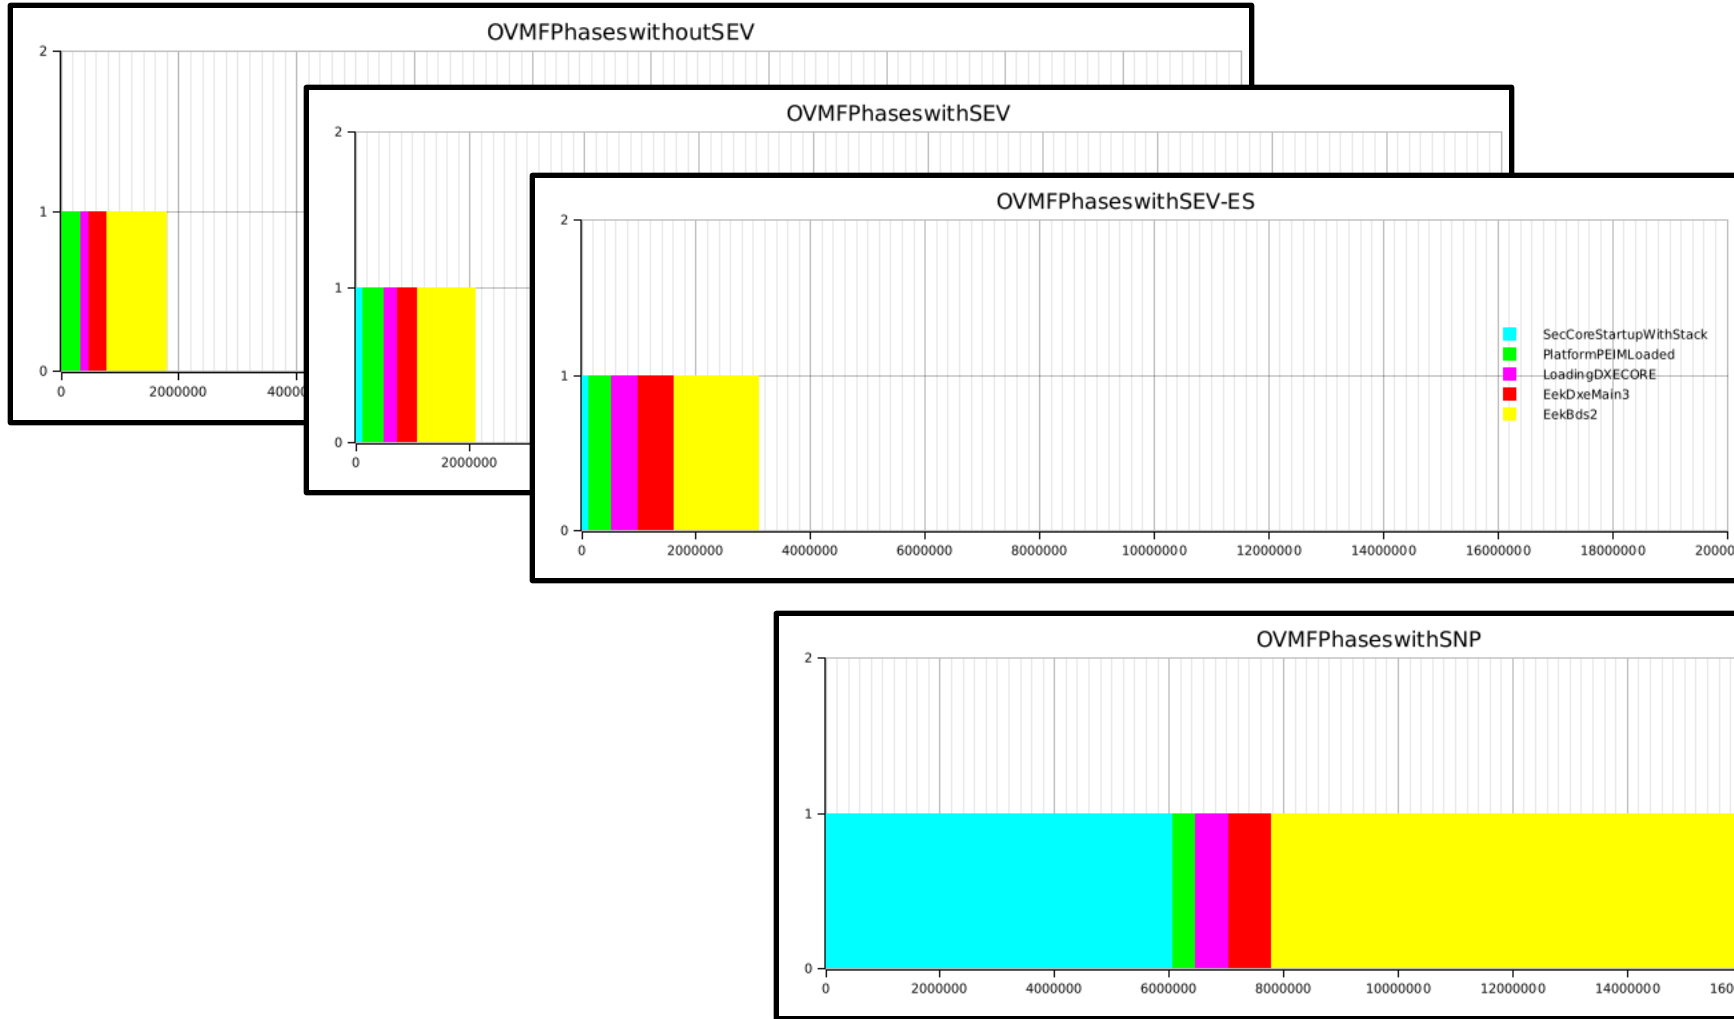
*Kernel based on 6.1.0-rc4*

- Neither of these papers is primarily focused on OVMF benchmarking
  - Holmes et al care about measurement
- There are lots of different guest configurations and software stacks



**Figure 4: Flame graph of the time spent booting a cVM.**

# Preliminary Results



*Performance Counter  
Frequency: 3,579,545 hz*

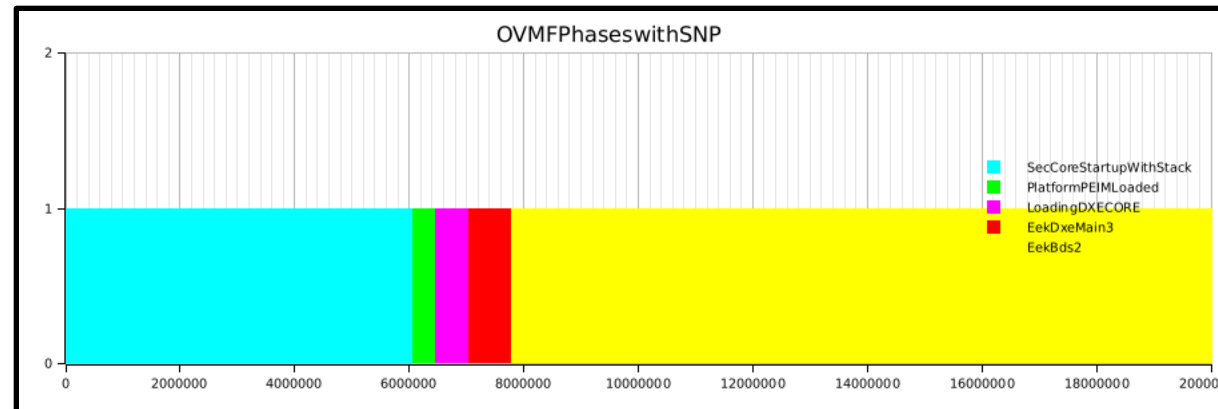
Host kernel is 6.8.0-rc5-next-20240221-snp-host-cc2568386ccb

# OVMF Debugging Techniques

- Debugger
- EFI Profiling
- Printing
  - Watch out for VMExits
- KVM Tracing
- EFI Shell
- Fuzzing
- Read the spec

# Memory Pre-Validation

- Three places
  - OvmfPkg/Sec/AmdSev.c
  - OvmfPkg/PlatformPei/AmdSev.c
  - OvmfPkg/AmdSevDxe/AmdSevDxe.c
- Plus
  - Unaccepted memory
  - And pre-encryption





# BDS Breakdown






```
//
// Do the platform init, can be customized by OEM/IBV
// Possible things that can be done in PlatformBootManagerBeforeConsole:
// > Update console variable: 1. include hot-plug devices; 2.
// > Register new Driver#### or Boot####
// > Register new Key####: e.g.: F12
// > Signal ReadyToLock event
// > Authentication action: 1. connect Auth devices; 2. Identify
//
PERF_INMODULE_BEGIN ("PlatformBootManagerBeforeConsole");
PlatformBootManagerBeforeConsole ();
PERF_INMODULE_END ("PlatformBootManagerBeforeConsole");
```

```
//
// Connect consoles
//
PERF_INMODULE_BEGIN ("EfiBootManagerConnectAllDefaultConsoles");
if (PcdGetBool (PcdConInConnectOnDemand)) {
    EfiBootManagerConnectConsoleVariable (ConOut);
    EfiBootManagerConnectConsoleVariable (ErrOut);
    //
    // Do not connect ConIn devices when lazy ConIn feature is ON.
    //
} else {
    EfiBootManagerConnectAllDefaultConsoles ();
}
PERF_INMODULE_END ("EfiBootManagerConnectAllDefaultConsoles");
```

# Audience Participation

- Which one is slowing us down?
  - device driver=virtio-net-pci,netdev=network-0,mac=ba:2f:08:16:18:aa,disable-modern=false,mq=on,vectors=4**
  - numa node,memdev=dimm1**
  - device virtio-scsi-pci,id=scsi,disable-modern=false**
  - device virtio-rng-pci,rng=rng0**



# Virtio-Rng??

- What does connecting consoles have to do with virtio-rng?
- Connecting consoles is surprisingly complex
  - See 3.15.3
  - Handles, Devices, Device Paths, Drivers
  - OVMF tries to bind most devices to most drivers
  - Could this be more enlightened?
- Why is Virtio-Rng slow?
  - It isn't
  - When it is enabled something else is slow
- Remove virtio-rng from Kata

# Something else?

```
33 R - - 0 1 7 PciRoot(0x0)
64 D - - 2 0 0 Primary Console Input Device
65 D - - 2 0 0 Primary Console Output Device
66 D - - 1 0 0 Primary Standard Error Device
86 D - - 1 0 0 PciRoot(0x0)/Pci(0x0,0x0)
87 B - - 1 1 1 QEMU Video PCI Adapter
88 D - - 1 0 0 PciRoot(0x0)/Pci(0x2,0x0)
89 D - - 1 3 0 PciRoot(0x0)/Pci(0x3,0x0)
8A B - - 1 1 3 PciRoot(0x0)/Pci(0x1F,0x0)
8B B - - 1 4 1 Sata Controller
8C D - - 1 0 0 PciRoot(0x0)/Pci(0x1F,0x3)
8E B - - 1 3 1 PciRoot(0x0)/Pci(0x1,0x0)/AcpiAdr(0x80010100)
91 B - - 1 1 1 PciRoot(0x0)/Pci(0x1F,0x0)/Serial(0x0)
92 D - - 1 0 0 PciRoot(0x0)/Pci(0x1F,0x0)/Serial(0x1)
93 B - - 1 3 1 PS/2 Keyboard Device
94 B - - 1 1 1 SIO Serial Port #0
95 B - - 1 5 3 VT-UTF8 Serial Console
96 D - - 1 2 0 QEMU QEMU DVD-ROM
```

**Without Virtio-Rng**

```
33 R - - 0 1 7 PciRoot(0x0)
64 D - - 2 0 0 Primary Console Input Device
65 D - - 2 0 0 Primary Console Output Device
66 D - - 1 0 0 Primary Standard Error Device
86 D - - 1 0 0 PciRoot(0x0)/Pci(0x0,0x0)
87 B - - 1 1 1 QEMU Video PCI Adapter
88 B - - 1 1 1 PciRoot(0x0)/Pci(0x2,0x0)
89 D - - 1 3 0 PciRoot(0x0)/Pci(0x3,0x0)
8A B - - 1 1 3 PciRoot(0x0)/Pci(0x1F,0x0)
8B B - - 1 4 1 Sata Controller
8C D - - 1 0 0 PciRoot(0x0)/Pci(0x1F,0x3)
8E B - - 1 3 1 PciRoot(0x0)/Pci(0x1,0x0)/AcpiAdr(0x80010100)
92 B - - 1 1 1 PciRoot(0x0)/Pci(0x1F,0x0)/Serial(0x0)
93 D - - 1 0 0 PciRoot(0x0)/Pci(0x1F,0x0)/Serial(0x1)
94 B - - 1 3 1 PS/2 Keyboard Device
95 B - - 1 1 1 SIO Serial Port #0
96 B - - 1 5 3 VT-UTF8 Serial Console
97 D - - 1 2 0 QEMU QEMU DVD-ROM
98 B - - 1 1 1 iPXE 825741 (0000:00:02.0, 52:54:00:12:34:56)
99 D - - 1 0 0 PciRoot(0x0)/Pci(0x2,0x0)/MAC(525400123456,0x1)/VenHw
4-1B9F-C54B-71E5-D6A16A5FB1AF)
```

**With Virtio-Rng**


# SCSI

- This is slow

```
//  
// Create the Maximum Attempts.  
//  
Status = IScsiCreateAttempts (PcdGet8 (PcdMaxIScsiAttemptNumber));  
if (EFI_ERROR (Status)) {  
    goto Error5;  
}
```

- Because it sets EFI variables
- On SEV(-ES) initializing QEMU Flash fails, so we avoid the slow path

# What about the new kernel?



- Writing to Pflash nvdata no longer slow
  - Perhaps because readonly memslots are no longer allowed for SVMs
  - Not the end of the story
- There is some overhead from pvalidate
  - It's less than expected

(6.11.0-rc5-snp-host-cc2568386)

# QEMU Flash

- Why are we emulating flash at all?
  - OVMF doesn't know how flash is provided
- QEMU allows potentially invalid configurations
- On SEV-ES we don't use QEMU flash, but it might be a bug

```
EFI_STATUS
QemuFlashWrite (
    IN      EFI_LBA  Lba,
    IN      UINTN    Offset,
    IN      UINTN    *NumBytes,
    IN      UINT8    *Buffer
)
{
    volatile UINT8  *Ptr;
    UINTN          Loop;

    ...

    // Restore flash to read mode
    //
    if (*NumBytes > 0) {
        QemuFlashPtrWrite (Ptr - 1, READ_ARRAY_CMD);
    }

    return EFI_SUCCESS;
}
```

## Example nvdata

```
Variable NV+RT+BS 'EFIGlobalVariable:BootOrder' DataSize = 0x04
Variable NV+RT+BS 'EFIGlobalVariable:Boot0001' DataSize = 0x58
Variable NV+RT+BS 'EFIGlobalVariable:ErrOut' DataSize = 0x49
Variable NV+RT+BS 'EFIGlobalVariable:ConIn' DataSize = 0x7A
Variable NV+RT+BS 'EFIGlobalVariable:ConOut' DataSize = 0x67
Variable NV+RT+BS 'EFIGlobalVariable:Key0001' DataSize = 0x0E
Variable NV+RT+BS 'EFIGlobalVariable:Key0000' DataSize = 0x0E
Variable NV+RT+BS 'EFIGlobalVariable:Lang' DataSize = 0x04
Variable NV+RT+BS 'EFIGlobalVariable:PlatformLang' DataSize = 0x03
Variable NV+RT+BS 'EFIGlobalVariable:Timeout' DataSize = 0x02
Variable NV+RT+BS 'EFIGlobalVariable:Boot0000' DataSize = 0x3E
Variable NV+RT+BS 'EFIGlobalVariable:Boot0000' DataSize = 0x3E
```

# Conclusion

- OVMF is not inherently slow with Confidential Computing
- But it is complex and difficult to analyze or optimize
  - Many configurations not regularly tested
- Let's figure out confidential warm starts
  
- Things to fix
  - Remove virtio-rng from Kata
  - Add QEMU warnings for invalid configurations
  - Take a close look at QEMU Flash and nvdata
  - Figure out SEC overhead