

FreeNAS - Bug #40476

Fix UEFI-CSM boot method for VMs

08/02/2018 05:37 AM - Chris Burge

Status:	Done		
Priority:	No priority		
Assignee:	Vaibhav Chauhan		
Category:	GUI (new)		
Target version:	11.2-BETA3		
Seen in:	11.2-BETA2	Needs Merging:	No
Severity:	Medium	Needs Automation:	No
Reason for Closing:		Support Suite Ticket:	n/a
Reason for Blocked:		Hardware Configuration:	
Needs QA:	No	ChangeLog Required:	No
Needs Doc:	No		

Description

Problem: The UEFI-CSM boot method remains in the VM guide with only a vague description of its applicability and use. VNC does not work with this boot method (VNC + uefi-csm selection currently not stopped in New UI – see <https://redmine.ixsystems.com/issues/40340> and old issue <https://redmine.ixsystems.com/issues/21150>).

This means you must use or generate an ISO that supports installation via a serial console. But a VM created with a cd-rom device with a such an ISO will not boot using the New UI.

Steps to re-create: Generate an iso that supports installation via serial console, for example using: <https://github.com/ffries/Debian-installer-with-serial-console> Use VM wizard to generate vm config. Remove VNC device, start and connect via “serial”. Error displayed in console:

```
[code]
Boot Failed. CDROM 0
Boot Failed. Harddisk 1
!!!! X64 Exception Type - 000000000000000D CPU Apic ID - 00000000 !!!!
RIP - 00000003FAF0FF5, CS - 0000000000000028, RFLAGS - 000000000010006
ExceptionData - 0000000000000000
RAX - 0000000000000000, RCX - 0000000000000008, RDX - 0000000000000050
RBX - 000000003F0C0618, RSP - 000000003FBDEA78, RBP - 000000003FBDEDD8
RSI - 00000000C0004120, RDI - 000000003FBDEA78
R8 - 0000000000000001, R9 - 000000003FBDEF7C, R10 - 00000000C0004120
R11 - 0000000000000000, R12 - 0000000000000800, R13 - 0000000000000000
R14 - 0000000000000000, R15 - 0000000000000000
DS - 0000000000000008, ES - 0000000000000008, FS - 0000000000000008
GS - 0000000000000008, SS - 0000000000000008
CR0 - 0000000080000023, CR2 - 0000000000000000, CR3 - 000000003FB7E000
CR4 - 0000000000000668, CR8 - 0000000000000000
DR0 - 0000000000000000, DR1 - 0000000000000000, DR2 - 0000000000000000
DR3 - 0000000000000000, DR6 - 00000000FFFF0FF0, DR7 - 0000000000004000
GDTR - 000000003FB68E98 000000000000003F, LDTR - 0000000000000000
IDTR - 000000003F6FA018 00000000000000FF, TR - 0000000000000000
FXSAVE_STATE - 000000003FBDE6D0
!!!! Find PE image
/wrkdirs/usr/ports/sysutils/uefi-edk2-bhyve-csm/work/uefi-edk2-0.1/Build/BhyveX64/RELEASE_GCC48/X64/UefiCpuPkg/CpuDxe/CpuDxe/DEBUG/CpuDxe.dll (ImageBase=000000003FAED000, EntryPoint=000000003FAED2AF) !!!!
[/code]
```

Messages in middlewared.log equates to this bhyve command:

```
bhyve -A -H -w -c 2 -m 1024
```

```
-s 0:0,hostbridge
-s 31,lpc
-l com1,/dev/nmdm4A -l bootrom,/usr/local/share/uefi-firmware/BHYVE_UEFI_CSM.fd
-s 3,virtio-net,tap0,mac=00:a0:98:65:dd:f4
-s 4,virtio-blk,/dev/zvol/NasPool2/VM/debcsm-kgcllv
-s 5,ahci-cd,/mnt/NasPool/home/chris/debian-9.4-amd64-CD-1.iso 4_debcsm
4_debcsm
```

Workaround: bhyve seems to be sensitive to which slots are used for which devices. I've not found anything in man bhyve or even on <https://lists.freebsd.org/pipermail/freebsd-virtualization/> that explains this, other than ahci devices need to be attached to slots 3,4,5 or 6, and nothing about how the boot order is decided. But the simple workaround is to change the slot order for the devices . At the CLI, this command boots the VM. allowing installation to take place:

```
bhyve -c 4 -m 1024M -A -H -w \
-s 0:0,hostbridge \
-s 31:0,lpc \
-s 4:0,ahci-cd,/mnt/NasPool/home/chris/debian-9.4-amd64-CD-1.iso \
-s 5:0,virtio-blk,/dev/zvol/NasPool2/VM/debcsm-kgcllv \
-s 3:0,virtio-net,tap20 \
-l com1,/dev/nmdm20A \
-l bootrom,/usr/local/share/uefi-firmware/BHYVE_UEFI_CSM.fd \
debcsm &
```

see image below debcsm1.jpeg

Once installed, it is possible to create a VM via the UI which will boot via UEFI-CSM.

middleware log shows:

```
[2018/07/17 10:57:16] (DEBUG) VMService.run():234 - Starting bhyve: bhyve -A -H -w -c 4 -m 1024 -s 0:0,hostbridge -s 31,lpc -l
com1,/dev/nmdm4A -l bootrom,/usr/local/share/uefi-firmware/BHYVE_UEFI_CSM.fd -s 3,virtio-net,tap0,mac=00:a0:98:65:dd:f4 -s
4,virtio-blk,/dev/zvol/NasPool2/VM/debcsm-kgcllv 4_debcsm
```

see images attached uefi-csm1.jpeg & uefi-csm2.jpeg

Possible solution: FreeNAS needs to investigate and provide a mechanism for selecting the boot order of devices attached to a VM. This also needed for VMs where multiple disk devices are attached. See this forum thread about how you cannot create a test VM of freenas with multiple disk attached via the UI:

<https://forums.freenas.org/index.php?threads/struggling-to-create-a-freenas-11-2-vm-in-iohype-bhyve.64005/#post-458658>

Subtasks:

Bug # 42371: Add order to VM devices API

Done

History

#1 - 08/02/2018 07:31 AM - Dru Lavigne

- Assignee changed from Release Council to Erin Clark
- Target version changed from Backlog to 11.2-BETA3

#3 - 08/02/2018 05:21 PM - Erin Clark

- Assignee changed from Erin Clark to Vaibhav Chauhan

#4 - 08/07/2018 12:04 PM - William Grzybowski

- Category changed from GUI (new) to Middleware
- Status changed from Unscreened to Blocked
- Assignee changed from Vaibhav Chauhan to William Grzybowski
- Severity changed from New to Medium
- Reason for Blocked set to Waiting for feedback

Chris, do you mean ahci-cd has to be first?

#5 - 08/12/2018 08:14 AM - Chris Burge

William Grzybowski wrote:

Chris, do you mean ahci-cd has to be first?

Yes, for the UEFI-CSM case I have outlined, in order for installation to take place.

In the case outlined in the Forum thread, the problem is after installation has taken place and additional HDD devices had been added to the VM. bhyve tries to boot from the wrong hdd device.

So how does bhyve pick the boot device? Is it by slot order? And how should this be reflected in the UI?

#6 - 08/13/2018 12:08 PM - William Grzybowski

- Status changed from *Blocked* to *Not Started*
- Reason for *Blocked* deleted (*Waiting for feedback*)

#7 - 08/20/2018 02:56 PM - William Grzybowski

- Category changed from *Middleware* to *GUI (new)*
- Assignee changed from *William Grzybowski* to *Vaibhav Chauhan*

Passing ticket to VB, middleware/backend work is in the subticket.

#8 - 08/22/2018 07:55 AM - Vaibhav Chauhan

- Status changed from *Not Started* to *In Progress*

in review: <https://github.com/freenas/webui/pull/1190>

#9 - 08/22/2018 12:14 PM - Dru Lavigne

- Subject changed from *Not possible to install a VM using UEFI-CSM boot method via the old/new UI. to Fix UEFI-CSM boot method for VMs*
- Status changed from *In Progress* to *Ready for Testing*
- Needs Doc changed from *Yes* to *No*
- Needs Merging changed from *Yes* to *No*

#10 - 09/07/2018 09:56 AM - Jeff Ervin

- File *VM UEFI CSM.png* added
- Status changed from *Ready for Testing* to *Passed Testing*
- Needs QA changed from *Yes* to *No*

Test Passes FreeNAS-11.2-MASTER-201809070900

#11 - 09/07/2018 10:03 AM - Dru Lavigne

- Status changed from *Passed Testing* to *Done*

Files

debcs1.jpeg	13.7 KB	08/02/2018	Chris Burge
uefi-csm-1.jpeg	30 KB	08/02/2018	Chris Burge
uefi-csm-2.jpeg	55.7 KB	08/02/2018	Chris Burge
VM UEFI CSM.png	217 KB	09/07/2018	Jeff Ervin