

DragonFlyBSD - Bug #3029

Running DflyBSD 4.8 on FreeBSD bhyve as a guest

04/15/2017 06:12 AM - iron

Status:	New	Start date:	04/14/2017
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	6.2		

Description

Hello,

I'd like to share my experience in running 4.8-RELEASE under FreeBSD 11-STABLE r316745M and report some bugs.

Install:

I used sysutils/vm-bhyve to create/manage VM. Here is config file:

```
graphics="yes"
graphics_listen="x.x.x.x"
graphics_port="5901"
graphics_res="800x600"
#graphics_wait="yes"
uefi="yes"
cpu=4
memory=8192M
network0_type="virtio-net"
network0_switch="public"
network1_type="virtio-net"
network1_switch="private"
#disk0_type="virtio-blk"
disk0_type="ahci-hd"
disk0_name="/dev/zvol/zroot/vm/php7-dfly_zvol"
disk0_dev="custom"
uuid="0ec14c29-2086-11e7-80ef-002590d96932"
network0_mac="00:50:56:00:C0:35"
network1_mac="58:9c:fc:09:2f:9c"
```

BHyve options:

```
Apr 15 16:00:34: [bhyve options: -c 6 -m 8192M -HwI bootrom,/vm/.config/BHYVE_UEFI.fd -U
0ec14c29-2086-11e7-80ef-002590d96932]
Apr 15 16:00:34: [bhyve devices: -s 0,hostbridge -s 31,lpc -s 4:0,virtio-blk,/dev/zvol/zroot/vm/php7-dfly_new_zvol -s
5:0,virtio-net,tap2,mac=00:50:56:00:C0:35 -s 5:1,virtio-net,tap3,mac=58:9c:fc:09:2f:9c -s
6:0,fbuf,tcp=148.251.46.105:5901,w=800,h=600,wait]
Apr 15 16:00:34: [bhyve console: -l com1,/dev/nmdm1A]
Apr 15 16:00:34: [bhyve iso device: -s 3:0,ahci-cd,/vm/.iso/dfly-x86_64-4.8.0_REL.iso]
```

In case of use virtio-blk for disk, installer can't format disk (see attached dfly1.png). If I click "Skip" (as I already have disk formatted) install process begins but it panics when trying to copy files (see attached dfly2.png and dfly3.png).

I tried to install use usual MBR boot record, but bhyve can't start VM with it. Finally, only with UEFI and disk as ahci-hd I could successfully install OS. After install I switched disk back to virtio-blk and here is what I saw:

```
-----
Copyright (c) 2003-2017 The DragonFly Project.
Copyright (c) 1992-2003 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
~ The Regents of the University of California. All rights reserved.
DragonFly v4.8.0-RELEASE #4: Sun Mar 26 20:55:22 EDT 2017
root@www.shiningsilence.com:/usr/obj/home/justin/release/4_8/sys/X86_64_GENERIC
i8254 clock: 1193393 Hzy.conf" 25L, 588C written
TSC invariant clock: 3499488936 Hz
CPU: Intel(R) Xeon(R) CPU E5-1650 v2 @ 3.50GHz (3499.49-MHz K8-class CPU)
Origin = "GenuineIntel" Id = 0x306e4 Stepping = 4
```

Features=0x9f83fbff<FPU,VME,DE,PSE,TSC,MSR,PAE,MCE,CX8,APIC,SEP,MTRR,PGE,MCA,CMOV,PAT,PSE36,MMX,FXSR,SS
E,SSE2,SS,HTT,PBE>
Features2=0xf69e6217<SSE3,PCLMULQDQ,DTES64,DS_CPL,SSSE3,CX16,xTPR,PCID,DCA,SSE4.1,SSE4.2,POPCNT,AESNI,X
SAVE,AVX,F16C,RDRND,VMM>
AMD Features=0x24100800<SYSCALL,NX,Page1GB,LM>
AMD Features2=0x1<LAHF>
Structured Extended Features=0x200<ENHMOVSB>
Thermal and PM Features=0x4<ARAT>
MONITOR/MWAIT Features=0x3<CST,INTBRK>
real memory = 8588931072 (8191 MB)
avail memory = 7974928384 (7605 MB)
lapic: divisor index 0, frequency 67097306 Hz
srat_probe: can't locate SRAT
SMI Frequency (worst case): 11235 Hz (89 us)
Initialize MI interrupts
TSC testing MP synchronization ...
TSC is MP synchronized
TSC: cputimer freq 109359029, shift 5
VMM: VMX is not supported by this Intel CPU
wdog: In-kernel automatic watchdog reset enabled
kbd1 at kbdmux0
md0: Malloc disk
interrupt uses mplock: swi_taskq
ACPI: RSDP 0x00000000BFB98014 000024 (v02 BHYVE)
ACPI: XSDT 0x00000000BFB970E8 00004C (v01 BHYVE BVFACP 00000001 01000013)
ACPI: FACP 0x00000000BFB96000 0000F4 (v04 BHYVE BVFACP 00000001 BHYV 00000001)
ACPI: DSDT 0x00000000BFB91000 0009A3 (v02 BHYVE BVSDT 00000001 INTL 20170303)
ACPI: FACS 0x00000000BFB9C000 000040
ACPI: HPET 0x00000000BFB95000 000038 (v01 BHYVE BVHPET 00000001 BHYV 00000001)
ACPI: APIC 0x00000000BFB94000 000062 (v01 BHYVE BVMADT 00000001 BHYV 00000001)
ACPI: MCFG 0x00000000BFB93000 00003C (v01 BHYVE BVMCFG 00000001 BHYV 00000001)
ACPI: SPCR 0x00000000BFB92000 000050 (v01 BHYVE BVSPCR 00000001 BHYV 00000001)
cryptosoft0: <software crypto> on motherboard
aesni0: <AES-CBC,AES-XTS> on motherboard
padlock0: No ACE support.
rdrand0: <RdRand RNG> on motherboard
acpi0: <BHYVE BVFACP> on motherboard
ACPI: 1 ACPI AML tables successfully acquired and loaded
ACPI FADT: SCI testing interrupt mode ...
ACPI FADT: SCI select level/low
acpi0: Power Button (fixed)
acpi_timer0 on acpi0
acpi_hpet0: <High Precision Event Timer> iomem 0xfed00000-0xfed003ff on acpi0
acpi_hpet0: frequency 10000000
pcib0: <ACPI Host-PCI bridge> port 0xcf8-0xcff,0x3c0-0x3df on acpi0
pci0: <ACPI PCI bus> on pcib0
pcib0: no PRT entry for 0.5.INTB
ahci0: <AHCI-PCI-SATA> mem 0xc0000000-0xc00003ff irq 5 at device 3.0 on pci0
ahci0: AE did not auto-clear!
ahci0: Intel hocus pocus
ahci0: AHCI 1.3 cap 0xff30ff25<S64A,NCQ,SSNTF,SMPS,SSS,SALP,SAL,SCLO,PMD,SSC,PSC,SXS> cap2 0x4<APST>, 6 ports,
32 tags/port, gen 3 (6Gbps)
ahci0.0: Found ATAPI "BHYVE SATA DVD ROM 001" serial="BHYVE-2C6F-F0BD-D7F8"
ahci0.0: tags=0/32 satacap=000e satacap2=0006 satafea=0020 NCQ=NO capacity=1.00MB
ahci0.0: f85=4218 f86=0000 f87=4000 WC=notsupp RA=notsupp SEC=notsupp
virtio_pci0: <VirtIO PCI Block adapter> port 0x2000-0x203f mem 0xc0002000-0xc0003fff irq 6 at device 4.0 on pci0
vtblk0: <VirtIO Block Adapter> on virtio_pci0
virtio_pci0: host features: 0x10000644 <RingIndirect,Topology,WriteCache,BlockSize,MaxNumSegs>
virtio_pci0: negotiated features: 0x10000244 <RingIndirect,WriteCache,BlockSize,MaxNumSegs>
virtio_pci1: <VirtIO PCI Network adapter> port 0x2040-0x205f mem 0xc0004000-0xc0005fff irq 7 at device 5.0 on pci0
vtnet0: <VirtIO Networking Adapter> on virtio_pci1
virtio_pci1: host features: 0x11018020 <RingIndirect,NotifyOnEmpty,Status,MrgRxBuf,MacAddress>
virtio_pci1: negotiated features: 0x11018020 <RingIndirect,NotifyOnEmpty,Status,MrgRxBuf,MacAddress>
vtnet0: MAC address: 00:50:56:00:c0:35
add dynamic link state
virtio_pci2: <VirtIO PCI Network adapter> port 0x2060-0x207f mem 0xc0006000-0xc0007fff irq 10 at device 5.1 on pci0
vtnet1: <VirtIO Networking Adapter> on virtio_pci2

```
virtio_pci2: host features: 0x11018020 <RingIndirect,NotifyOnEmpty,Status,MrgRxBuf,MacAddress>
virtio_pci2: negotiated features: 0x11018020 <RingIndirect,NotifyOnEmpty,Status,MrgRxBuf,MacAddress>
vtnet1: MAC address: 58:9c:fc:09:2f:9c
add dynamic link state
isab0: <PCI-ISA bridge> at device 31.0 on pci0
isa0: <ISA bus> on isab0
atkbd0: <Keyboard controller (i8042)> port 0x64,0x60 irq 1 on acpi0
atkbd0: <AT Keyboard> irq 1 on atkbd0
kbd0 at atkbd0
psm0: <PS/2 Mouse> irq 12 on atkbd0
interrupt uses mlock: psm0
psm0: model Generic PS/2 mouse, device ID 0
sio0: <16550A-compatible COM port> port 0x3f8-0x3ff irq 4 on acpi0
sio0: type 16550A, console
sio1: <16550A-compatible COM port> port 0x2f8-0x2ff irq 3 on acpi0
sio1: type 16550A
pmtimer0 on isa0
vga0: <Generic ISA VGA> at port 0x3c0-0x3df iomem 0xa0000-0xbfff on isa0
device_probe_and_attach: vga0 attach returned 6
hpt27xx: no controller detected.
CAM: Configuring 1 busses
CAM: finished configuring all busses
cd0 at ahci0 bus 0 target 0 lun 0
cd0: <BHYVE BHYVE DVD-ROM 001> Removable CD-ROM SCSI-0 device
cd0: 600.000MB/s transfers
cd0: cd present [0 x 2048 byte records]
no B_DEVMAGIC (bootdev=0)
Mounting root from hammer:serno/BHYVE-4A50-7F6A-B929.s1d
no disk named 'serno/BHYVE-4A50-7F6A-B929.s1d'
hammer_vfs_mount: can't find devvp
hammer_vfs_mount: No root volume found!
Root mount failed: 22
```

```
Manual root filesystem specification:
<fstype>:<device> Specify root (e.g. ufs:da0s1a)
?           List valid disk boot devices
panic      Just panic
abort      Abort manual input
```

```
mountroot> panic: Root mount failed, startup aborted.
cpuid = 0
Trace beginning at frame 0xffffffff817f4f10
panic() at panic+0x236 0xffffffff805e5d76
panic() at panic+0x236 0xffffffff805e5d76
vfs_mountroot() at vfs_mountroot+0xfe 0xffffffff8065e00e
mi_startup() at mi_startup+0x84 0xffffffff805b2004
Debugger("panic")

CPU0 stopping CPUs: 0x0000000e
stopped
Stopped at Debugger+0x7c: movb $0,0xd9cb79(%rip)
db>
```

I tried to change serno/BHYVE-4A50-7F6A-B929.s1d mention in /boot/loadenv.conf, /etc/rc.conf and /etc/fstab to vbd0s1d or set "mountroot> hammer:vbd0s1d", but every time kernel get panics.

From time to time VM freezes after: TSC testing MP synchronization ...
...and I need to reset VM a few times until it pass ahead this stage.
If I run "truss -p <bhyve PID>" while it freezed I see repeating messages like:

```
ioctl(3,0xc0907601 { IORW 0x76('v'), 1, 144 },0xddbebe20) = 0 (0x0)
ioctl(3,0xc0907601 { IORW 0x76('v'), 1, 144 },0xdd9eae20) = 0 (0x0)
ioctl(3,0xc0907601 { IORW 0x76('v'), 1, 144 },0xddbebe20) = 0 (0x0)
ioctl(3,0xc0907601 { IORW 0x76('v'), 1, 144 },0xdd9eae20) = 0 (0x0)
ioctl(3,0xc0907601 { IORW 0x76('v'), 1, 144 },0xddbebe20) = 0 (0x0)
ioctl(3,0xc0907601 { IORW 0x76('v'), 1, 144 },0xdd7e9e20) = 0 (0x0)
```

```
ioctl(3,0xc0907601 { IORW 0x76('v'), 1, 144 },0xddbebe20) = 0 (0x0)
ioctl(3,0xc0907601 { IORW 0x76('v'), 1, 144 },0xdd9eae20) = 0 (0x0)
ioctl(3,0xc0907601 { IORW 0x76('v'), 1, 144 },0xdd5e8e20) = 0 (0x0)
ioctl(3,0xc0907601 { IORW 0x76('v'), 1, 144 },0xdd9eae20) = 0 (0x0)
ioctl(3,0xc0907601 { IORW 0x76('v'), 1, 144 },0xdd7e9e20) = 0 (0x0)
ioctl(3,0xc0907601 { IORW 0x76('v'), 1, 144 },0xddbebe20) = 0 (0x0)
ioctl(3,0xc0907601 { IORW 0x76('v'), 1, 144 },0xdd9eae20) = 0 (0x0)
ioctl(3,0xc0907601 { IORW 0x76('v'), 1, 144 },0xddbebe20) = 0 (0x0)
ioctl(3,0xc0907601 { IORW 0x76('v'), 1, 144 },0xdd9eae20) = 0 (0x0)
ioctl(3,0xc0907601 { IORW 0x76('v'), 1, 144 },0xddbebe20) = 0 (0x0)
```

From time to time disk subsystem freezes and in dmesg appear timeouts from ahci:

```
sched_itd: stray interrupt 7 on cpu0
sched_itd: stray interrupt 7 on cpu0
ahci1.0: CMD TIMEOUT state=5 slot=10
glb-status 0x00000000
cmd-reg 0x1ca17<CPS,CR,FR,FRE,POD,SUD,ST>
port_status 0x0
sactive=00000400 active=00000000 expired=00000000
sact=00000000 ci=00000000
STS=50
ahci1.0: disk_rw: timeout
sched_itd: stray interrupt 7 on cpu0
sched_itd: stray interrupt 7 on cpu0
sched_itd: stray interrupt 7 on cpu0
sched_itd: stray interrupt 7 on cpu0
sched_itd: stray interrupt 7 on cpu0
sched_itd: stray interrupt 7 on cpu0
sched_itd: stray interrupt 7 on cpu0
sched_itd: 10 stray interrupts 7 on cpu0 - there will be no further reports
ahci1.0: CMD TIMEOUT state=5 slot=8
glb-status 0x00000000
cmd-reg 0x1c817<CPS,CR,FR,FRE,POD,SUD,ST>
port_status 0x0
sactive=00000100 active=00000000 expired=00000000
sact=00000000 ci=00000000
STS=50
ahci1.0: disk_rw: timeout
ahci1.0: CMD TIMEOUT state=5 slot=18
glb-status 0x00000000
cmd-reg 0x1d217<CPS,CR,FR,FRE,POD,SUD,ST>
port_status 0x0
sactive=00040000 active=00000000 expired=00000000
sact=00000000 ci=00000000
STS=50
ahci1.0: disk_rw: timeout
-----
```

Currently only three major bugs:

1. installer fails to install OS on virtio-blk disk
2. virtio-blk driver doesn't work.
3. ahci-hd timeouts from time to time.

I have tried to install v4.9.0.176.g1ee016 and faced the same problems.

I don't think that the problems with disk could be BHyve/FreeBSD related because all other OSes as guests work fine (FreeBSD, Linux, Windows).

Sorry for long post. I hope all the info above would be useful for somebody for further investigations.

History

#1 - 04/26/2017 10:27 AM - asomers

See <https://github.com/churchers/vm-bhyve/pull/156> for a working vm-bhyve template for DragonFly 4.8

#2 - 04/27/2017 10:32 AM - iron

asomers wrote:

See <https://github.com/churchers/vm-bhyve/pull/156> for a working vm-bhyve template for DragonFly 4.8

It was mine bug report which is mainly about bug in documentation. I already solved problem with EFI boot and described configuration here (forgot to update report there).

In this bug report I described the problems which make DflyBSD 4.8 unusable under bhyve.

Files

dfly1.png	27 KB	04/15/2017	iron
dfly2.png	23.7 KB	04/15/2017	iron
dfly3.png	25.9 KB	04/15/2017	iron