
Subject: debuginfo for kernel

Posted by [amorsen](#) on Wed, 09 Jan 2008 22:47:00 GMT

[View Forum Message](#) <> [Reply to Message](#)

I am trying to nail down a performance problem with routing.

CPU: AMD64 processors, speed 2194.5 MHz (estimated)

Counted CPU_CLK_UNHALTED events (Cycles outside of halt state) with a unit mask of 0x00
(No unit mask) count 100000

vma	samples	%	linenr	info	symbol name
ffffffff811e83a2	586492	25.2725	(no location information)		dev_hard_start_xmit
ffffffff811e83a2	68	0.0116	(no location information)		
ffffffff811e83a4	63	0.0107	(no location information)		
ffffffff811e83aa	61	0.0104	(no location information)		
ffffffff811e83b1	92	0.0157	(no location information)		
ffffffff811e83c5	55	0.0094	(no location information)		
ffffffff811e83cb	3	5.1e-04	(no location information)		
ffffffff811e83d2	61	0.0104	(no location information)		
ffffffff811e83d4	1	1.7e-04	(no location information)		
ffffffff811e83dd	75	0.0128	(no location information)		
ffffffff811e83ee	2	3.4e-04	(no location information)		
ffffffff811e83f2	492753	84.0170	(no location information)		

[..]

So a lot of the kernel time is spent in the dev_hard_start_xmit function, and in particular on the address ffffffff811e83f2. The big question is -- what is that instruction doing?

I think I need the debuginfo package to find out, but I can't find the debuginfo package.

/Benny

Subject: Re: debuginfo for kernel

Posted by [dev](#) on Thu, 10 Jan 2008 07:14:07 GMT

[View Forum Message](#) <> [Reply to Message](#)

1. you can recompile kernel yourself from sources to get all the debuginfo required
2. or do objdump -dr vmlinux-your-vesrion and find the instruction bu this address in assembler dump.

it's really strange and should not be one instruction which takes ~85% cpu time.

What performance problem do you try to investigate?

Subject: Re: debuginfo for kernel
Posted by [amorsen](#) on Thu, 10 Jan 2008 18:39:01 GMT
[View Forum Message](#) <> [Reply to Message](#)

I tried with gdb, but it is a bit difficult to figure out which bit in the assembly corresponds to which bit of the source code.

```
(gdb) disassemble dev_hard_start_xmit
Dump of assembler code for function dev_hard_start_xmit:
0xffffffff811e83a2 <dev_hard_start_xmit+0>:  push  %r13
0xffffffff811e83a4 <dev_hard_start_xmit+2>:  push  %r12
0xffffffff811e83a6 <dev_hard_start_xmit+4>:  push  %rbp
0xffffffff811e83a7 <dev_hard_start_xmit+5>:  mov   %rdi,%rbp
0xffffffff811e83aa <dev_hard_start_xmit+8>:  push  %rbx
0xffffffff811e83ab <dev_hard_start_xmit+9>:  push  %r13
0xffffffff811e83ad <dev_hard_start_xmit+11>: cmpq  $0x0,(%rdi)
0xffffffff811e83b1 <dev_hard_start_xmit+15>:  mov   %rsi,%r13
0xffffffff811e83b4 <dev_hard_start_xmit+18>:  jne   0xffffffff811e85ad
<dev_hard_start_xmit+523>
0xffffffff811e83ba <dev_hard_start_xmit+24>:  cmpq  $0xffffffff8169b160,4926875(%rip)  #
0xffffffff8169b160 <ptype_all>
0xffffffff811e83c5 <dev_hard_start_xmit+35>:  je    0xffffffff811e84ed
<dev_hard_start_xmit+331>
0xffffffff811e83cb <dev_hard_start_xmit+41>:  cmpl  $0x0,6474374(%rip)  #
0xffffffff81814e58 <netstamp_needed>
0xffffffff811e83d2 <dev_hard_start_xmit+48>:  je    0xffffffff811e83df <dev_hard_start_xmit+61>
0xffffffff811e83d4 <dev_hard_start_xmit+50>:  callq 0xffffffff81045b0c <ktime_get_real>
0xffffffff811e83d9 <dev_hard_start_xmit+55>:  mov   %rax,0x18(%rbp)
0xffffffff811e83dd <dev_hard_start_xmit+59>:  jmp   0xffffffff811e83e7
<dev_hard_start_xmit+69>
0xffffffff811e83df <dev_hard_start_xmit+61>:  movq  $0x0,0x18(%rdi)
0xffffffff811e83e7 <dev_hard_start_xmit+69>:  mov   4926834(%rip),%rdx  #
0xffffffff8169b160 <ptype_all>
0xffffffff811e83ee <dev_hard_start_xmit+76>:  lea  0xffffffffffffd0(%rdx),%r12
0xffffffff811e83f2 <dev_hard_start_xmit+80>:  mov  0x30(%r12),%rax
0xffffffff811e83f7 <dev_hard_start_xmit+85>:  prefetch0 (%rax)
0xffffffff811e83fa <dev_hard_start_xmit+88>:  cmp  $0xffffffff8169b160,%rdx
0xffffffff811e8401 <dev_hard_start_xmit+95>:  je    0xffffffff811e84ed
<dev_hard_start_xmit+331>
[.]
```

I guess I will have to build the kernel myself. The problematic instruction is at dev_hard_start_xmit+80.

The performance problem is dropped packets when routing more than about 8000 packets/s.

Subject: Re: debuginfo for kernel

Posted by [amorsen](#) on Sat, 12 Jan 2008 23:32:16 GMT

[View Forum Message](#) <> [Reply to Message](#)

I tried rebuilding the kernel rpm to get debuginfo-packages. It didn't work.

I set builddebuginfo to 1, but no debuginfo packages were generated from rpmbuild -ba kernel-ovz.spec. I'm building for x86_64.
