
Subject: localhost issues on multi-VM machine

Posted by [copperfrog](#) on Wed, 16 Jan 2013 17:47:49 GMT

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

Hi,

I have an OpenVZ host machine which runs 10 guests on it. The system is configured so that each of those VMs is using a certain amount of RAM and CPU's without over-allocating the physical machine. I.e. it's a 32 core machine with 64 GB RAM, all machines together are configured to use 63GB of RAM and 28 CPU cores.

I run Apache on all of the VMs and also Munin. In order to monitor correct behavior and patterns Munin uses 'apache_accesses', 'apache_processes' and 'apache_volume' on each of these VMs. Now the nature of Munin is to come along every 5 minutes and run all those scripts and retrieve the values. I have observed that if these scripts try to check /server-status via localhost (http : //localhost:80/server-status?auto) they time out on some of the VMs. This can be also observed by running a stupid while loop requesting the status with curl from the command line:

```
while [ 1 ]; do `curl http : //localhost/server-status?auto > /dev/null`; done
```

It will run until Munin comes along and then curl will report 'host unreachable' for about 30 seconds until it recovers.

I checked the beancounters and monitor the lo interface and do not see any errors on either of them.

It looks to me like the localhost interface is shared across all VMs and is running out of capacity to serve all of the requests. The math would dictate that at 10 VMs with 3 scripts running each there are 30 requests within a very small time frame. Maybe I'm wrong but I'd expect that load to be acceptable.

I can alleviate the problem by using the actual IP address of the VM (assigned to venet0) for these calls. So the venet0 interface seems to have a much higher capacity for requests. Another way to relieve the affected VMs is by turning off Munin for some of the other VMs; I can run up to 7 VMs fine; 8 and more cause time outs.

Does anybody have an idea what the problem might be?

Here some of the information requested by the stickys:

- Host running CentOS 6.3
- 8 VMs running CentOS 6.3 (2 affected)
- 1 VM running Ubuntu 10.04 (fine)
- 1 VM running CentOS 5.8 (fine)

One of the affected VMs:

>ifconfig

```
lo      Link encap:Local Loopback
        inet addr:127.0.0.1  Mask:255.0.0.0
        inet6 addr: ::1/128 Scope:Host
        UP LOOPBACK RUNNING  MTU:16436  Metric:1
        RX packets:1845251 errors:0 dropped:0 overruns:0 frame:0
        TX packets:1845251 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
```

RX bytes:189093532 (180.3 MiB) TX bytes:189093532 (180.3 MiB)

venet0 Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
inet addr:127.0.0.1 P-t-P:127.0.0.1 Bcast:0.0.0.0 Mask:255.255.255.255
UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1
RX packets:13190677 errors:0 dropped:0 overruns:0 frame:0
TX packets:17361191 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:5631175102 (5.2 GiB) TX bytes:17001995532 (15.8 GiB)

venet0:0 Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
inet addr:10.10.20.45 P-t-P:10.10.20.45 Bcast:10.10.20.45 Mask:255.255.255.255
UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1

> ip rule list

0: from all lookup local

32766: from all lookup main

32767: from all lookup default

> ip route list table all

169.254.0.0/16 dev venet0 scope link metric 1002

default dev venet0 scope link

broadcast 127.255.255.255 dev lo table local proto kernel scope link src 127.0.0.1

local 10.10.20.45 dev venet0 table local proto kernel scope host src 10.10.20.45

broadcast 10.10.20.45 dev venet0 table local proto kernel scope link src 10.10.20.45

broadcast 127.0.0.0 dev lo table local proto kernel scope link src 127.0.0.1

local 127.0.0.1 dev lo table local proto kernel scope host src 127.0.0.1

local 127.0.0.1 dev venet0 table local proto kernel scope host src 127.0.0.1

local 127.0.0.0/8 dev lo table local proto kernel scope host src 127.0.0.1

unreachable ::/96 dev lo metric 1024 error -101 mtu 16436 advmss 16376 hoplimit 4294967295

unreachable ::ffff:0.0.0.0/96 dev lo metric 1024 error -101 mtu 16436 advmss 16376 hoplimit 4294967295

unreachable 2002:a00::/24 dev lo metric 1024 error -101 mtu 16436 advmss 16376 hoplimit 4294967295

unreachable 2002:7f00::/24 dev lo metric 1024 error -101 mtu 16436 advmss 16376 hoplimit 4294967295

unreachable 2002:a9fe::/32 dev lo metric 1024 error -101 mtu 16436 advmss 16376 hoplimit 4294967295

unreachable 2002:ac10::/28 dev lo metric 1024 error -101 mtu 16436 advmss 16376 hoplimit 4294967295

unreachable 2002:c0a8::/32 dev lo metric 1024 error -101 mtu 16436 advmss 16376 hoplimit 4294967295

unreachable 2002:e000::/19 dev lo metric 1024 error -101 mtu 16436 advmss 16376 hoplimit 4294967295

unreachable 3ffe:ffff::/32 dev lo metric 1024 error -101 mtu 16436 advmss 16376 hoplimit 4294967295

default dev venet0 metric 1 mtu 1500 advmss 1440 hoplimit 4294967295

unreachable default dev lo table unspec proto kernel metric -1 error -101 hoplimit 255

```
local ::1 via :: dev lo table local proto none metric 0 mtu 16436 advmss 16376 hoplimit 4294967295
unreachable default dev lo table unspec proto kernel metric -1 error -101 hoplimit 255
```

```
> iptables -L
```

```
Chain INPUT (policy ACCEPT)
```

```
target prot opt source destination
```

```
Chain FORWARD (policy ACCEPT)
```

```
target prot opt source destination
```

```
Chain OUTPUT (policy ACCEPT)
```

```
target prot opt source destination
```

```
> tcpdump -i lo -e host 127.0.0.1
```

```
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
```

```
listening on lo, link-type EN10MB (Ethernet), capture size 65535 bytes
```

```
09:19:47.896467 00:00:00:00:00:00 (oui Ethernet) > 00:00:00:00:00:00 (oui Ethernet), ethertype IPv4 (0x0800), length 74: localhost.localdomain.33038 > localhost.localdomain.http: Flags [S], seq 2650559739, win 32768, options [mss 16396,sackOK,TS val 3647429171 ecr 3647428160,nop,wscale 9], length 0
```

```
09:19:47.896502 00:00:00:00:00:00 (oui Ethernet) > 00:00:00:00:00:00 (oui Ethernet), ethertype IPv4 (0x0800), length 74: localhost.localdomain.http > localhost.localdomain.33038: Flags [S.], seq 3952410189, ack 2650559740, win 32768, options [mss 16396,sackOK,TS val 3647429171 ecr 3647429171,nop,wscale 9], length 0
```

```
09:19:47.896530 00:00:00:00:00:00 (oui Ethernet) > 00:00:00:00:00:00 (oui Ethernet), ethertype IPv4 (0x0800), length 66: localhost.localdomain.33038 > localhost.localdomain.http: Flags [.], ack 1, win 64, options [nop,nop,TS val 3647429171 ecr 3647429171], length 0
```

```
09:19:47.896626 00:00:00:00:00:00 (oui Ethernet) > 00:00:00:00:00:00 (oui Ethernet), ethertype IPv4 (0x0800), length 248: localhost.localdomain.33038 > localhost.localdomain.http: Flags [P.], seq 1:183, ack 1, win 64, options [nop,nop,TS val 3647429171 ecr 3647429171], length 182
```

```
09:19:47.896664 00:00:00:00:00:00 (oui Ethernet) > 00:00:00:00:00:00 (oui Ethernet), ethertype IPv4 (0x0800), length 66: localhost.localdomain.http > localhost.localdomain.33038: Flags [.], ack 183, win 67, options [nop,nop,TS val 3647429171 ecr 3647429171], length 0
```

```
09:19:47.897343 00:00:00:00:00:00 (oui Ethernet) > 00:00:00:00:00:00 (oui Ethernet), ethertype IPv4 (0x0800), length 428: localhost.localdomain.http > localhost.localdomain.33038: Flags [P.], seq 1:363, ack 183, win 67, options [nop,nop,TS val 3647429172 ecr 3647429171], length 362
```

```
09:19:47.897369 00:00:00:00:00:00 (oui Ethernet) > 00:00:00:00:00:00 (oui Ethernet), ethertype IPv4 (0x0800), length 66: localhost.localdomain.33038 > localhost.localdomain.http: Flags [.], ack 363, win 67, options [nop,nop,TS val 3647429172 ecr 3647429172], length 0
```

```
09:19:47.897589 00:00:00:00:00:00 (oui Ethernet) > 00:00:00:00:00:00 (oui Ethernet), ethertype IPv4 (0x0800), length 66: localhost.localdomain.33038 > localhost.localdomain.http: Flags [F.], seq 183, ack 363, win 67, options [nop,nop,TS val 3647429172 ecr 3647429172], length 0
```

```
09:19:47.897697 00:00:00:00:00:00 (oui Ethernet) > 00:00:00:00:00:00 (oui Ethernet), ethertype IPv4 (0x0800), length 66: localhost.localdomain.http > localhost.localdomain.33038: Flags [F.], seq 363, ack 184, win 67, options [nop,nop,TS val 3647429172 ecr 3647429172], length 0
```

```
09:19:47.897724 00:00:00:00:00:00 (oui Ethernet) > 00:00:00:00:00:00 (oui Ethernet), ethertype IPv4 (0x0800), length 66: localhost.localdomain.33038 > localhost.localdomain.http: Flags [.], ack
```

```
364, win 67, options [nop,nop,TS val 3647429172 ecr 3647429172], length 0
09:19:48.907586 00:00:00:00:00:00 (oui Ethernet) > 00:00:00:00:00:00 (oui Ethernet), ethertype
IPv4 (0x0800), length 74: localhost.localdomain.33040 > localhost.localdomain.http: Flags [S], seq
1138244790, win 32768, options [mss 16396,sackOK,TS val 3647430182 ecr
3647429172,nop,wscale 9], length 0
09:19:48.907615 00:00:00:00:00:00 (oui Ethernet) > 00:00:00:00:00:00 (oui Ethernet), ethertype
IPv4 (0x0800), length 74: localhost.localdomain.http > localhost.localdomain.33040: Flags [S.],
seq 1886956057, ack 1138244791, win 32768, options [mss 16396,sackOK,TS val 3647430182
ecr 3647430182,nop,wscale 9], length 0
09:19:48.907632 00:00:00:00:00:00 (oui Ethernet) > 00:00:00:00:00:00 (oui Ethernet), ethertype
IPv4 (0x0800), length 66: localhost.localdomain.33040 > localhost.localdomain.http: Flags [.], ack
1, win 64, options [nop,nop,TS val 3647430182 ecr 3647430182], length 0
09:19:48.907693 00:00:00:00:00:00 (oui Ethernet) > 00:00:00:00:00:00 (oui Ethernet), ethertype
IPv4 (0x0800), length 248: localhost.localdomain.33040 > localhost.localdomain.http: Flags [P.],
seq 1:183, ack 1, win 64, options [nop,nop,TS val 3647430182 ecr 3647430182], length 182
09:19:48.907715 00:00:00:00:00:00 (oui Ethernet) > 00:00:00:00:00:00 (oui Ethernet), ethertype
IPv4 (0x0800), length 66: localhost.localdomain.http > localhost.localdomain.33040: Flags [.], ack
183, win 67, options [nop,nop,TS val 3647430182 ecr 3647430182], length 0
09:19:48.908267 00:00:00:00:00:00 (oui Ethernet) > 00:00:00:00:00:00 (oui Ethernet), ethertype
IPv4 (0x0800), length 430: localhost.localdomain.http > localhost.localdomain.33040: Flags [P.],
seq 1:365, ack 183, win 67, options [nop,nop,TS val 3647430183 ecr 3647430182], length 364
09:19:48.908306 00:00:00:00:00:00 (oui Ethernet) > 00:00:00:00:00:00 (oui Ethernet), ethertype
IPv4 (0x0800), length 66: localhost.localdomain.33040 > localhost.localdomain.http: Flags [.], ack
365, win 67, options [nop,nop,TS val 3647430183 ecr 3647430183], length 0
09:19:48.908705 00:00:00:00:00:00 (oui Ethernet) > 00:00:00:00:00:00 (oui Ethernet), ethertype
IPv4 (0x0800), length 66: localhost.localdomain.33040 > localhost.localdomain.http: Flags [F.],
seq 183, ack 365, win 67, options [nop,nop,TS val 3647430183 ecr 3647430183], length 0
09:19:48.908803 00:00:00:00:00:00 (oui Ethernet) > 00:00:00:00:00:00 (oui Ethernet), ethertype
IPv4 (0x0800), length 66: localhost.localdomain.http > localhost.localdomain.33040: Flags [F.],
seq 365, ack 184, win 67, options [nop,nop,TS val 3647430183 ecr 3647430183], length 0
09:19:48.908830 00:00:00:00:00:00 (oui Ethernet) > 00:00:00:00:00:00 (oui Ethernet), ethertype
IPv4 (0x0800), length 66: localhost.localdomain.33040 > localhost.localdomain.http: Flags [.], ack
366, win 67, options [nop,nop,TS val 3647430183 ecr 3647430183], length 0
```

^C

20 packets captured

40 packets received by filter

0 packets dropped by kernel

Host:

> uname -a

Linux vmh21.copperfroghosting.net 2.6.32-042stab065.3 #1 SMP Mon Nov 12 21:59:14 MSK
2012 x86_64 x86_64 x86_64 GNU/Linux

>ifconfig

bond0 Link encap:Ethernet HWaddr BC:30:5B:EE:2B:BC
inet addr:10.10.20.181 Bcast:10.10.20.255 Mask:255.255.255.0

inet6 addr: fe80::be30:5bff:feee:2bbc/64 Scope:Link
UP BROADCAST RUNNING MASTER MULTICAST MTU:1500 Metric:1
RX packets:726264499 errors:0 dropped:0 overruns:0 frame:0
TX packets:747256747 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:511078335328 (475.9 GiB) TX bytes:624997159521 (582.0 GiB)

eth0 Link encap:Ethernet HWaddr BC:30:5B:EE:2B:BC
UP BROADCAST RUNNING SLAVE MULTICAST MTU:1500 Metric:1
RX packets:558321526 errors:0 dropped:0 overruns:0 frame:0
TX packets:259011771 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:488649254639 (455.0 GiB) TX bytes:211547180801 (197.0 GiB)
Memory:dcb00000-dcc00000

eth1 Link encap:Ethernet HWaddr BC:30:5B:EE:2B:BD
UP BROADCAST RUNNING SLAVE MULTICAST MTU:1500 Metric:1
RX packets:167942973 errors:0 dropped:0 overruns:0 frame:0
TX packets:488244976 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:22429080689 (20.8 GiB) TX bytes:413449978720 (385.0 GiB)
Memory:dcc00000-dcd00000

lo Link encap:Local Loopback
inet addr:127.0.0.1 Mask:255.0.0.0
inet6 addr: ::1/128 Scope:Host
UP LOOPBACK RUNNING MTU:16436 Metric:1
RX packets:475169 errors:0 dropped:0 overruns:0 frame:0
TX packets:475169 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:97271832 (92.7 MiB) TX bytes:97271832 (92.7 MiB)

venet0 Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
inet6 addr: fe80::1/128 Scope:Link
UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1
RX packets:696720709 errors:0 dropped:0 overruns:0 frame:0
TX packets:661464289 errors:0 dropped:37 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:611456529535 (569.4 GiB) TX bytes:496828263061 (462.7 GiB)

> iptables -L
Chain INPUT (policy ACCEPT)
target prot opt source destination

Chain FORWARD (policy ACCEPT)
target prot opt source destination

Chain OUTPUT (policy ACCEPT)

target	prot	opt	source	destination
--------	------	-----	--------	-------------