Subject: Bridging inside the CT, snort in-line?!
Posted by vitorallo on Sat, 31 Jan 2009 11:28:28 GMT
View Forum Message <> Reply to Message

Hello Folks,

nice to be part of the community, I just joined.
Few words about openvz. It's awesome. I started using it only one week ago and it already changed my life!!!! it is a great project and a tremendous contribution to the virtualization world....

I liked it so much that I decided to start a small project using it. Unfortunately I stuck with an architectural problem.

I read already that bridge-utils and bridging inside the ct is not possible. This makes perfectly sense. It is plenty of possibilities to bridge outside at the host machine level.

However, I am trying to install snort in-line inside the ct. I am so crazy to think about IPS inside CT. No problem with recompiling it; I just used a debootstrap minimal machine. The real problem is to let the traffic flowing into the VM, bridging inside and allowing iptable queue to grab it, in order for snort to process the analysis. Therefore, bridging cannot be out of the ct. Even working with a stick into the ct, after the learning phase the bridge will send the traffic directly out to the second interface (without taking care of the stick)

Let me explain. This what I would like to do:

eth0(host)--(bridge or...)--veth101.0----eth0(CT)
 ..... .!
 ..... .! bridge
 ..... .! snort-inline
 ..... .!
eth1(host)--(bridge2 or..)---veth101.1---eth1(CT)

Not that complicated architecture, but snort runs inside and needs to lookup the packets from the ct inner eb or iptables.

I looked around the internet but I did not find anything useful. Every efforts in bridging for openvz is done outside. Maybe I was not very good in searching around. I tried also some horizontal solution like using iptables to copy packets but I don't want to add so much overhead (affecting performances).

Right now, I am looking for something, modules and so on, something to compile as external pluggable module into the ovz kernel ... uhmm ... it wouldn't be very clean.

I don't think nobody tried to do something similar. So, I am posting here to grab your suggestions, idea or testing. I am pretty new and I might miss the way to do it help
I would like to go ahead with openvz - I don't want to switch on xen!!! I will be happy to share trough this forum any other further experience of mine .

talk to you soon guys,

vito

Subject: Re: Bridging inside the CT, snort in-line?!
Posted by maratrus on Mon, 02 Feb 2009 09:15:57 GMT
View Forum Message <> Reply to Message

Hello,

I've never used snort in-line so my questions might look silly but I really don't understand why do you need bridge interface inside CT. Could you possibly explain in more detail why it's so important?

As far as I understand your scheme the bridge on the HN unites physical eth* and virtual veth* interfaces. And what do you want to unite into bridge inside CT?

Quote:
The real problem is to let the traffic flowing into the VM, bridging inside and allowing iptable queue to grab it

Why it is so important if a packet passed bridge interface?

Thank you.

Subject: Re: Bridging inside the CT, snort in-line?!
Posted by vitorallo on Mon, 02 Feb 2009 14:47:17 GMT
View Forum Message <> Reply to Message

Of course it is not a silly question.
Deploying snort in line means create an IPS and not an IDS. IPSes - intrusion detection and prevention - work by analyzing the traffic and reacting "in line" by applying a response like: block, rewrite and so on. To do that they need to stay in the middle of the traffic flow.

Any good IPS (I work for Internet Security Systems) has two interface. IPS are deployed literally in the middle of the traffic trunk.

internet ----- firewall ---- (portA)IPS(portB)---- inner network

usually IPSes are transparent! there is a bridge/special nic driver that brings the traffic over the two interfaces and blocks it when needed

Snort can work as an IPS, it is actually an IDS. To do that, it needs to stay upon a network bridge.

it can grab traffic from libpcap or the iptables QUEUE. For the IPS/Bridge we need the second feature.

Using little bit of fancy ASCII art...

..iptables QUEUE.. <---------- snort
!--------------------!
!----linux bridge---!
+.....................+
A---- TRAFFIC --- B

My aim is to scale security architecture, creating an easy re-distributable template with snort in line to secure the
Virtual Infrastructure!!! THAT WILL BE COOL... of course I
want to share it open.

Unfortunately, to do it I need to bridge inside the CT, I need
linux bridge.ko (bridge utils) active inside and iptables running.

No problem for iptables, snort and all the rest.. but I cannot bridge eth0 and eth1 (virtual interfaces) inside the CT...

Imagine how cool it would be....to have something like

----Host networking----VirtualWorld---IPS transparent---multiple OpenVZ machines

---

Subject: Re: Bridging inside the CT, snort in-line?!
Posted by maratrus on Mon, 02 Feb 2009 17:32:36 GMT
View Forum Message <> Reply to Message

First of all, thank you for clarifications. So if I'm right, the main desire is not running two instances of snort inside CT. Instead it would be great to join interfaces and run only one instance of snort. But running two instances of snort is a solution anyway (and it's no so bad). What do you think?

Unfortunately it's impossible to create/delete bridges inside CT because they are not virtualized. There is a bug http://bugzilla.openvz.org/show_bug.cgi?id=831 (and you could describe the necessity of this feature from new point of view).

May be it worth trying the following (ugly) workaround: using only one eth0 interface inside CT and two ip addresses on it (alias) and appropriate routing rules on the HN.
To be more plain:

```
 eth0     eth1
+-|--------|-+
| |   HN   | |
|            |
|     veth   |
```

```
|  +---|---+ |
|  | eth0  | |
|  | eth0:0| |
|  |  VE   | |
|  +-------+ |
+------------+
```

On the HN the following routing rules:

```
ip r add $CT_IP_1 dev $VETH
ip r add $CT_IP_2 dev $VETH
```

Subject: Re: Bridging inside the CT (within VPS)
Posted by vitorallo on Mon, 09 Feb 2009 22:46:49 GMT
View Forum Message <> Reply to Message

Hello...

first, sorry for being so late in answering: my wife gave birth of my wonderful daughter, so I have been busy with a different business.

how can express my gratitude for your help. Yes the solution with iproute is something I thought at the beginning. Indeed, it is a solution and it is based on layer3 configuration. as you said, it is not really clean. It is actually something different from my original thought.

However, you showed me the bug then I followed the thread about "bridging withing vps". Devel provided a solution for an old old 2.6 kernel. With the help of a friend of mine, he went into the kernel code and he patched the 2.6.18 to do the same stuff. YESS he is awesome!!!

Bridge started into the VM. Tomorrow we will test the solution and check if it works properly. We are aware that this impacts the security profile..

I am available to follow up on the bug and to provide information if they want to work it in the mainstream.. I suppose also my friend is full available...

I will keep this thread updated, getting it useful to somebody else...

Subject: Re: Bridging inside the CT, snort in-line?!
Posted by ivani on Wed, 06 Apr 2011 13:56:37 GMT
View Forum Message <> Reply to Message

Hi vitorallo,

I'm looking for a solution for my problem with the snort IDS.
The parent host run openvz, and I've installed the CentOS 5.5, this is output of uname:

```
uname -a
Linux snortlab 2.6.18-194.8.1.el5.028stab070.5 #1 SMP Fri Sep 17 19:10:36 MSD 2010 i686 i686
i386 GNU/Linux
```

I'm not sure what kind of interface is venet0:0, I thought it was xen.

I tried this:

```
snort -vv -i lo
Running in packet dump mode

        --== Initializing Snort ==--
Initializing Output Plugins!
pcap DAQ configured to passive.
Acquiring network traffic from "lo".
Decoding Ethernet

        --== Initialization Complete ==--

   ,,_     -*> Snort! <*-
  o"  )~   Version 2.9.0.4 IPv6 GRE (Build 110)
   ''''    By Martin Roesch & The Snort Team:
           Copyright (C) 1998-2011 Sourcefire, Inc., et al.
           Using libpcap version 1.1.1
           Using PCRE version: 6.6 06-Feb-2006
           Using ZLIB version: 1.2.3

Commencing packet processing (pid=21572
```

Well, this works fine. But, if I try:

```
snort -vv -i venet0:0
Running in packet dump mode

        --== Initializing Snort ==--
Initializing Output Plugins!
pcap DAQ configured to passive.
Acquiring network traffic from "venet0:0".
Decoding Linux SLL

        --== Initialization Complete ==--

   ,,_     -*> Snort! <*-
```

```
o"  )~   Version 2.9.0.4 IPv6 GRE (Build 110)
 ""    By Martin Roesch & The Snort Team:
          Copyright (C) 1998-2011 Sourcefire, Inc., et al.
          Using libpcap version 1.1.1
          Using PCRE version: 6.6 06-Feb-2006
          Using ZLIB version: 1.2.3

Commencing packet processing (pid=5776)
Can't acquire (-1) - cooked-mode frame doesn't have room for sll header!
```

And the snort can't start.

I've googled many pages, forums, mail lists, but I'm still lost about this weird problem.

Any ideas?

Thank you so much.

Regards,

Ivani Nascimento
