

---

Subject: IPCop template

Posted by [marc1911](#) on Thu, 02 Aug 2007 18:34:41 GMT

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

---

Hi there,

I'm very impressed by OpenVZ and I'm thinking of consolidating my Linux OS's which are on VMWare to OpenVZ. This because of the lack of speed I'm experiencing. One of the things I need to overcome is the migration of IPCop v1.4.16 to OpenVZ. I'ven't found an IPCop image between the OS images on Openvz.org. Wouldn't it be great if IPCop would be available as a template? Are there any people around that fee the same or are there perhaps any already engaged in developing such a template? I would be really interested!

Regards,

Marc

---

---

Subject: Re: IPCop template

Posted by [ugob](#) on Thu, 02 Aug 2007 18:42:08 GMT

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

---

I think it may not be worth the effort. IpCop, as a firewall, would be better installed by itself, unless you are testing, as the default VE in OpenVZ uses venet, which is a virtual interface, maybe veth would let you get something close but what you want, but I'm not sure.

[http://wiki.openvz.org/Virtual\\_Ethernet\\_device](http://wiki.openvz.org/Virtual_Ethernet_device)

[http://wiki.openvz.org/Differences\\_between\\_venet\\_and\\_veth](http://wiki.openvz.org/Differences_between_venet_and_veth)

---

---

Subject: Re: IPCop template

Posted by [marc1911](#) on Thu, 02 Aug 2007 19:01:25 GMT

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

---

I'm working with IPCop for sometime now on VMWare Server without any problems. Works like a charm! I don't see why a firewall couldn't be in a virtual environment. Do you foresee any security issues when running IPCop as a virtual computer?

Furthermore I do not understand your remark about the default VE and the virtual networkinterface. Are you saying that there maybe not enough services to tackle the networkinterfaces on IPCop? I could use more hardware networkinterfaces if there's a problem on the virtual side?

---

---

Subject: Re: IPCop template  
Posted by [ugob](#) on Thu, 02 Aug 2007 19:23:10 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

I don't think you easily assign physical interface to a VM in openvz like you can in VmWare, that is what I meant.

---

Subject: Re: IPCop template  
Posted by [marc1911](#) on Thu, 02 Aug 2007 21:44:41 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Okay. But what about the section "Moving Network Adapter to Virtual Private Server" on page 68 of the Chapter "Advanced Tasks" in the OpenVZ Users Guide? I've read that just after your reply and it says that there's a possibility for a VPS to directly access the physical network adapter. The following command would make that possible:

```
# vzctl set 101 --netdev_add eth1 --save
```

It has some limitations though, one of them is that the device is only accessible to the VPS. Hardware node and other VPS's thus excluded.

---

Subject: Re: IPCop template  
Posted by [ugob](#) on Thu, 02 Aug 2007 21:48:41 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Ok, I was just warning you, I can't help you more than that because I only used venet up to now. I think it would be better to make sure that this is working before trying to make a firewall template.

Regards,

Ugo

---

Subject: Re: IPCop template  
Posted by [marc1911](#) on Thu, 02 Aug 2007 21:56:30 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Ugo,

That's ofcourse good advice! I shall test it before I put my mind to construct a firewall template. But than again I was also hoping that there would be anybody around in the OpenVZ community who already did it or something similar. It can't be a new idea!

---

---

Subject: Re: IPCop template

Posted by [nschembr](#) on Tue, 30 Oct 2007 19:01:20 GMT

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

---

I'm going to move a UML, User Mode Linux image to openVZ. If I have time, I will try to create the template for Smoothwall, ipcop.

Smoothwall is on the 2.6 kernel. ipcop 1.5. is slated to be the first 2.6 kernel release.

Any one else working on something like this, ie. VE controlled firewall with gui and fun stuff for users to play with.

---

### File Attachments

1) [openvz.smoothwall.jpg](#), downloaded 618 times

2) [docs.network.openvz\\_.pdf](#), downloaded 449 times

---

---

Subject: Re: IPCop template

Posted by [nschembr](#) on Fri, 02 Nov 2007 04:02:35 GMT

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

---

Ok, the smooth wall port looks good so far.

I created br0 on 10.200.50.202.

I created a debian ve from the template. I tested the bridge. Everything works great. I can get from the ve to the public internet and from the public internet to the ve.

I removed the debian ve from the bridge.

I created a tar ball from the Physical to VE wiki. step by step converted Smoothwall 3.0 to a ve.

I used the Virtual Ethernet device wiki to create 2 real interfaces in the smoothwall VE.

I started smooth wall and played with setup until I found the red and green interfaces.

At this point smooth wall has one interface marked red, 10.200.50.213, and one interface marked green , 10.200.1.1

From smooth wall I can ping google.com and 10.200.50.1 (the modem)

I can ping 10.200.50.213 from the Hardware Node and other hosts on the local net.

I can start the debian ve. this time it is connected to the green bridge. Using dhcp it pulls an address from smooth wall. debian now has 10.200.1.96. and a gateway of 10.200.1.1.

Debian can ping 10.200.1.1, 10.200.50.213, 10.200.50.202 but not 10.200.50.1 or 10.200.50.152

Is this an issue with smoothwall and nat?

why would nat not start inside the ve? what should I look for?

Time to read more of the wiki.

ps. It took less then an hour to do the port. Great job making openvz easy to use.

## File Attachments

1) [test1 diag.pdf](#), downloaded 645 times

---

---

Subject: Re: IPCop template

Posted by [nschembr](#) on Fri, 02 Nov 2007 05:02:02 GMT

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

---

It looks like a strange bridge configuration. I'm going to build the bridge by hand without ip  
Addresses to see what I can get.

but it is now time for sleep.

---

---

Subject: Re: IPCop template

Posted by [nschembr](#) on Fri, 02 Nov 2007 21:22:04 GMT

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

---

Help. Ok, I'm stuck.

Hardware node# ipconfig

```
77g-br  Link encap:Ethernet HWaddr 00:18:51:8C:B1:31
        inet addr:10.200.1.20 Bcast:10.200.1.255 Mask:255.255.255.0
        inet6 addr: fe80::218:51ff:fe8c:b131/64 Scope:Link
        UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
        RX packets:2 errors:0 dropped:0 overruns:0 frame:0
        TX packets:11 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:56 (56.0 b) TX bytes:902 (902.0 b)
```

```
br0     Link encap:Ethernet HWaddr 00:18:51:D4:35:40
        inet addr:192.168.1.202 Bcast:192.168.1.255 Mask:255.255.255.0
        inet6 addr: fe80::21c:c0ff:fe05:a51/64 Scope:Link
        UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
        RX packets:16831 errors:0 dropped:0 overruns:0 frame:0
        TX packets:9324 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:24588860 (23.4 MiB) TX bytes:660477 (644.9 KiB)
```

```
eth1   Link encap:Ethernet HWaddr 00:1C:C0:05:0A:51
inet6 addr: fe80::21c:c0ff:fe05:a51/64 Scope:Link
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:16830 errors:0 dropped:0 overruns:0 frame:0
TX packets:10288 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:100
RX bytes:24892265 (23.7 MiB) TX bytes:794884 (776.2 KiB)
Base address:0x30e0 Memory:e0300000-e0320000

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:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)

venet0 Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)

veth101.0 Link encap:Ethernet HWaddr 00:18:51:D4:35:40
inet6 addr: fe80::218:51ff:fed4:3540/64 Scope:Link
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:46 errors:0 dropped:0 overruns:0 frame:0
TX packets:55 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:4735 (4.6 KiB) TX bytes:4837 (4.7 KiB)

veth104.0 Link encap:Ethernet HWaddr 00:18:51:8C:B1:31
inet6 addr: fe80::218:51ff:fe8c:b131/64 Scope:Link
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:996 errors:0 dropped:0 overruns:0 frame:0
TX packets:170 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:83173 (81.2 KiB) TX bytes:11963 (11.6 KiB)

veth105.0 Link encap:Ethernet HWaddr 00:18:51:DE:9A:89
inet6 addr: fe80::218:51ff:fede:9a89/64 Scope:Link
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:936 errors:0 dropped:0 overruns:0 frame:0
TX packets:60 errors:0 dropped:0 overruns:0 carrier:0
```

collisions:0 txqueuelen:0  
RX bytes:74816 (73.0 KiB) TX bytes:3512 (3.4 KiB)

veth105.1 Link encap:Ethernet HWaddr 00:18:51:CF:F1:7D  
inet6 addr: fe80::218:51ff:fecf:f17d/64 Scope:Link  
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1  
RX packets:154 errors:0 dropped:0 overruns:0 frame:0  
TX packets:988 errors:0 dropped:0 overruns:0 carrier:0  
collisions:0 txqueuelen:0  
RX bytes:10067 (9.8 KiB) TX bytes:80629 (78.7 KiB)

hardware node#route

Kernel IP routing table

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
10.200.1.0	*	255.255.255.0	U	0	0	0	77g-br
192.168.1.0	*	255.255.255.0	U	0	0	0	br0
default	192.168.1.1	0.0.0.0	UG	0	0	0	br0

Hardware node# brctl show

bridge name	bridge id	STP enabled	interfaces
77g-br	8000.0018518cb131	no	veth105.1 veth104.0
br0	8000.001851d43540	no	eth1 veth101.0 veth105.0

Smoothwall

fw77 (root) / \$ ifconfig

eth0 Link encap:Ethernet HWaddr 00:18:51:EB:A2:3F  
inet addr:192.168.1.212 Bcast:192.168.1.255 Mask:255.255.255.0  
inet6 addr: fe80::218:51ff:feeb:a23f/64 Scope:Link  
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1  
RX packets:75 errors:0 dropped:0 overruns:0 frame:0  
TX packets:1338 errors:0 dropped:0 overruns:0 carrier:0  
collisions:0 txqueuelen:0  
RX bytes:4166 (4.0 Kb) TX bytes:107856 (105.3 Kb)

eth1 Link encap:Ethernet HWaddr 00:18:51:55:11:34  
inet addr:10.200.1.1 Bcast:10.200.1.255 Mask:255.255.255.0  
inet6 addr: fe80::218:51ff:fe55:1134/64 Scope:Link  
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1  
RX packets:1386 errors:0 dropped:0 overruns:0 frame:0  
TX packets:163 errors:0 dropped:0 overruns:0 carrier:0  
collisions:0 txqueuelen:0  
RX bytes:113557 (110.8 Kb) TX bytes:10319 (10.0 Kb)

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:33 errors:0 dropped:0 overruns:0 frame:0
TX packets:33 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:3016 (2.9 Kb) TX bytes:3016 (2.9 Kb)
```

RX bytes:3016 (2.9 Kb) TX bytes:3016 (2.9 Kb)

```
fw77 (root) / $ route
```

```
Kernel IP routing table
```

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
192.168.1.0	*	255.255.255.0	U 0	0	0	eth0	
10.200.1.0	*	255.255.255.0	U 0	0	0	eth1	
default	192.168.1.1	0.0.0.0	UG 0	0	0	eth0	

```
fw77 (root) / $
```

The fw can access 192.168.1.1. and the internet.

The VE host

```
root@localhost:/# ifconfig
```

```
eth0 Link encap:Ethernet HWaddr 00:18:51:27:7F:0A
inet addr:10.200.1.25 Bcast:10.200.1.255 Mask:255.255.255.0
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:185 errors:0 dropped:0 overruns:0 frame:0
TX packets:1602 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:12383 (12.0 KiB) TX bytes:133237 (130.1 KiB)
```

```
root@localhost:/# route
```

```
Kernel IP routing table
```

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
10.200.1.0	*	255.255.255.0	U 0	0	0	eth0	
default	fw77	0.0.0.0	UG 0	0	0	eth0	

```
root@localhost:/#
```

The ve host can access 192.168.1.202, 192.168.1.212 , and 10.200.1.1. but not 192.168.1.1

If I run tcpdump on the hardware node and i ping 192.168.1.1, I can see the traffic from host 10.200.1.25. This sounds like the nat on smoothwall is not working.

Diag test1 diag.pdf is true. I have changed 10.200.50.0/24 to 192.169.1.0/24 for testing.

How do you I turn on Nat inside the ve?

---

---

Subject: Re: IPCop template

Posted by [nschembr](#) on Mon, 05 Nov 2007 01:04:01 GMT

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

---

Ok, I'm stuck.

The port is very easy, but Ip table inside the ve is an issue.

tt's very fast to make the smoothwall template.

I used [http://wiki.openvz.org/Physical\\_to\\_VE](http://wiki.openvz.org/Physical_to_VE) to create the image.

I used the Virtual ethernet devices can be joined in one bridge section from [http://wiki.openvz.org/Veth#Simple\\_configuration\\_with\\_virtual\\_ethernet\\_device](http://wiki.openvz.org/Veth#Simple_configuration_with_virtual_ethernet_device) to create two bridges. one green and one red. The HN shared the red bridge with the ve.

The VE has a issue with iptables. I worked with "Re: iptables with nat inside guest", <http://forum.openvz.org/index.php?t=msg&goto=22599&>

An Iptables expert is needed for this port.

I'm going to install kvm and run the guest inside the vm. I will post the howto

---

---

Subject: Re: IPCop template

Posted by [sjdean](#) on Fri, 30 May 2008 18:46:58 GMT

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

---

I have three servers

- 1) Smoothwall
- 2) Fedora - DNS/DHCP/Mail/Web
- 3) Fedora - Asterisk

I plan to merge them altogether to save electricity, but I always fancied keeping my Smoothwall. I like the idea of OpenVZ. How easy is it all?

How did you get on with your Smoothwall setup under OpenVZ?

ta  
Simon

---