

Re: [solved] Using racoon–negotiated IPsec with ipfw and natd

Source: <http://www.derkeiler.com/Mailing-Lists/FreeBSD-Security/2003-11/0002.html>

From: Mark Johnston (*mjohnston_at_skyweb.ca*)

Date: 11/03/03

Date: Mon, 3 Nov 2003 11:05:32 -0600

To: cjclark@alum.mit.edu

"Crist J. Clark" <crstjc@comcast.net> wrote:

> For packets entering the system from the network, the processing
> order is,
>
> (network) ----> ipfw ----> IPsec ----> (remainder of IP stack)
>
> And outgoing,
>
> (system) ----> IPsec ----> ipfw ----> (network)
>
> (It's actually a bit more hairy that that, incoming IPsec processed
> packets actually get reinjected into the stack below ipfw processing,
> but skip ipfw on the second pass, unless IPSEC_FILTERGIF is set.)
> Notice I didn't explicitly say where natd(8) happens because ipfw(8)
> passes packets to natd(8) and that is completely under your control.
>
> The problem is that the addresses on the packets has been rewritten
> before they are being set out the external interface where IPsec
> processing would happen.

Perfect! Thank you! That's exactly the explanation I needed.

> Ouch. Mixing bridging, NAT, and IPsec. (I should talk, my bastion host
> at home has one interface with my coax cable connection, another to my
> NATed LAN, another to my NATed WLAN which also is all tunneled through
> IPsec or PPTP since WEP is broken, and finally some PPP dial-up
> interfaces to call into the office. No bridging there, though! Only
> bridge on test boxes on the internal LAN.)
>
> I don't understand is what breaks if you just do,
>
> 500 divert natd ip from 192.168.15.0/24 to any out via fxp0
> 600 divert natd ip from any to me in via fxp0
>
> And lose 700. Is there a reason to NAT stuff between the internal

FreeBSD–Security: Re: [solved] Using racoon–negotiated IPsec with ipfw and natd

> *network and DMZ?*

There is – I'm not the DMZ's gateway, and NAT means not having to add static routes to all the DMZ boxes. The legacy box that this is replacing (a 3Com SuperStack 3000–series firewall appliance) actually allowed packets to DMZ hosts by responding to ARPs for their IPs on the WAN port, then invisibly proxying packets to them, and doing the same trick for DMZ–WAN traffic. I avoided that fate by bridging, but reconfiguring the remote box that actually is the DMZ gateway wasn't an option.

For the archives:

Dynamic (roaming) IPsec was not working with racoon on a firewall also running natd. The problem was that natd was rewriting the packets as they came in, because of an ipfw rule matching on the internal interface, and by the time the packets made it to the IPsec layer, they no longer matched the SP. It was fixed by changing the ipfw rule to match only outgoing packets, which will already have been processed by IPsec by the time they get to ipfw on the trip out.

Specifically, I've made one tiny change to my ruleset. I replaced this rule:

```
00500 divert 8669 ip from 192.168.15.0/24 to not me recv txp0
```

with this:

```
00500 divert 8669 ip from 192.168.15.0/24 to not me out recv txp0
```

Adding "out" prevents ipfw from diverting the packet to natd on its way in. On the way out, the packet has been through IPsec and will no longer match the "from 192.168.15.0/24" criterion, saving it from diversion again. If you don't also use a DMZ with bridging, you can do it a lot more easily, as Crist describes above.

Thanks a lot for your help,
Mark

freebsd–security@freebsd.org mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd–security>

To unsubscribe, send any mail to "freebsd–security–unsubscribe@freebsd.org"