Subject: simple internet router in VE?
Posted by digidax on Tue, 20 Jan 2009 09:07:23 GMT
View Forum Message <> Reply to Message

Hello,

I'm trying to setup a simple router with NAT to get internet access from IP 172.16.0.99 at the LAN (eth1) to the WAN (eth0).

interfaces on HN:
ifconfig
eth0 Link encap:Ethernet Hardware Adresse 00:15:17:4B:0A:50
inet Adresse:192.168.130.208 Bcast:192.168.130.255 Maske:255.255.255.0
inet6 Adresse: fe80::215:17ff:fe4b:a50/64 GÃ¼ltigkeitsbereich:Verbindung
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:4194 errors:0 dropped:0 overruns:0 frame:0
TX packets:7392 errors:0 dropped:0 overruns:0 carrier:0
Kollisionen:0 SendewarteschlangenlÃ¤nge:1000
RX bytes:490557 (479.0 KiB) TX bytes:1777986 (1.6 MiB)
Speicher:feb80000-feba0000 javascript://

eth1 Link encap:Ethernet Hardware Adresse 00:15:17:4B:0A:51
inet Adresse:172.16.0.3 Bcast:172.16.0.255 Maske:255.255.255.0
inet6 Adresse: fe80::215:17ff:fe4b:a51/64 GÃ¼ltigkeitsbereich:Verbindung
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:406 errors:0 dropped:0 overruns:0 frame:0
TX packets:180 errors:0 dropped:0 overruns:0 carrier:0
Kollisionen:0 SendewarteschlangenlÃ¤nge:100
RX bytes:31650 (30.9 KiB) TX bytes:61399 (59.9 KiB)
Speicher:febe0000-fec00000

lo Link encap:Lokale Schleife
inet Adresse:127.0.0.1 Maske:255.0.0.0
inet6 Adresse: ::1/128 GÃ¼ltigkeitsbereich:Maschine
UP LOOPBACK RUNNING MTU:16436 Metric:1
RX packets:8 errors:0 dropped:0 overruns:0 frame:0
TX packets:8 errors:0 dropped:0 overruns:0 carrier:0
Kollisionen:0 SendewarteschlangenlÃ¤nge:0
RX bytes:560 (560.0 b) TX bytes:560 (560.0 b)

venet0 Link encap:UNSPEC Hardware Adresse
00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
UP BROADCAST PUNKTZUPUNKT RUNNING NOARP MTU:1500 Metric:1
RX packets:6623 errors:0 dropped:0 overruns:0 frame:0
TX packets:1535 errors:0 dropped:0 overruns:0 carrier:0
Kollisionen:0 SendewarteschlangenlÃ¤nge:0
RX bytes:1659375 (1.5 MiB) TX bytes:136607 (133.4 KiB)

but
ip r l
172.16.0.4 dev venet0 scope link
192.168.130.209 dev venet0 scope link
192.168.130.0/24 dev eth0 proto kernel scope link src 192.168.130.208
172.16.0.0/24 dev eth1 proto kernel scope link src 172.16.0.3
169.254.0.0/16 dev eth1 scope link
default via 192.168.130.254 dev eth0

asks me, why is network 169.254.0.0/16 present?

From the VE I can ping all host at the WAN, also DNS resolving works. I can also ping from the VE to all Clients located at LAN. From LAN I can ping eth1 and can open a website from apache inside VE on port 80. DNS resolving from internet hosts works also inside LAN because a BIND is also running inside VE.

on VE I'm using:

iptables -A FORWARD -s 172.16.0.99 -j ACCEPT
iptables -A POSTROUTING -t nat -s 172.16.0.99 -j MASQUERADE

I have run a tcpdump while the client tryed: http://www.google.de

on HN: # tcpdump host 172.16.0.99
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on eth0, link-type EN10MB (Ethernet), capture size 96 bytes
08:54:10.189894 IP 172.16.0.99.kiosk > bw-in-f99.google.com.http: S
4160233486:4160233486(0) win 65535 <mss 1460,nop,nop,sackOK>
08:54:13.127257 IP 172.16.0.99.kiosk > bw-in-f99.google.com.http: S
4160233486:4160233486(0) win 65535 <mss 1460,nop,nop,sackOK>
08:54:19.136236 IP 172.16.0.99.kiosk > bw-in-f99.google.com.http: S
4160233486:4160233486(0) win 65535 <mss 1460,nop,nop,sackOK>
08:54:31.156345 IP 172.16.0.99.veracity > bw-in-f103.google.com.http: S
3111893581:3111893581(0) win 65535 <mss 1460,nop,nop,sackOK>
08:54:34.160112 IP 172.16.0.99.veracity > bw-in-f103.google.com.http: S
3111893581:3111893581(0) win 65535 <mss 1460,nop,nop,sackOK>
08:54:40.068256 IP 172.16.0.99.veracity > bw-in-f103.google.com.http: S
3111893581:3111893581(0) win 65535 <mss 1460,nop,nop,sackOK>
08:54:40.770259 IP 172.16.0.99.nfs > ipcop.localdomain.ntp: NTPv3, Client, length 48
08:54:52.090269 IP 172.16.0.99.kyoceranetdev > bw-in-f104.google.com.http: S
3303672974:3303672974(0) win 65535 <mss 1460,nop,nop,sackOK>
08:54:55.093254 IP 172.16.0.99.kyoceranetdev > bw-in-f104.google.com.http: S
3303672974:3303672974(0) win 65535 <mss 1460,nop,nop,sackOK>
08:55:01.102152 IP 172.16.0.99.kyoceranetdev > bw-in-f104.google.com.http: S
3303672974:3303672974(0) win 65535 <mss 1460,nop,nop,sackOK>

08:55:13.121419 IP 172.16.0.99.jstel > bw-in-f147.google.com.http: S 1832507979:1832507979(0) win 65535 <mss 1460,nop,nop,sackOK>

13 packets captured
26 packets received by filter
0 packets dropped by kernel

on VE: # tcpdump host 172.16.0.99
tcpdump: WARNING: arptype 65535 not supported by libpcap - falling back to cooked socket
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on venet0, link-type LINUX_SLL (Linux cooked), capture size 96 bytes
00:59:29.992900 IP 172.16.0.99.dls-monitor > 172.16.0.4.domain: 26096+ A? www.google.de. (31)
00:59:30.277110 IP 172.16.0.4.domain > 172.16.0.99.dls-monitor: 26096 6/7/0 CNAME[|domain]

2 packets captured
6 packets received by filter
0 packets dropped by kernel

Where is the routing problem from HN to VE?