
Subject: ****FIXED**** DNS not working in VE
Posted by [jchamilton](#) on Wed, 12 Jul 2006 21:51:34 GMT
[View Forum Message](#) <> [Reply to Message](#)

I've searched the forum, wiki, and mailing lists and the only thread I could find in which the issue is "solved" is here: <http://forum.openvz.org/index.php?t=tree&th=761&mid=3947&&rev=&reveal=>

The problem is I cannot get a response for DNS queries executed from my VE's if iptables is running on the HN.

I have ip_conntrack enabled in /etc/modules.conf

I have a nameserver config'd in /etc/resolv.conf

nsswitch.conf has "hosts: files dns"

I can connect to the VE via ssh. (using IP addr)

I can connect to other machines on the network from the VE via ssh. (using IP addr)

If iptables is stopped, DNS lookups on the VE work. (using ping, dig, and getent)

If iptables is running dig gives the error: "connection timed out; no servers could be reached" even if I specify the nameserver on the command line.

If iptables is running ping just says: "unknown host ..."

I'm pretty sure at one point it was working though and that's what really mystifies me...

I've also restarted vz and iptables on the HN with no joy.

Here's what the firewall tables look like:

Table: nat

Chain PREROUTING (policy ACCEPT)

```
target  prot opt source          destination
```

Chain POSTROUTING (policy ACCEPT)

```
target  prot opt source          destination
```

Chain OUTPUT (policy ACCEPT)

```
target  prot opt source          destination
```

Table: mangle

Chain PREROUTING (policy ACCEPT)

```
target  prot opt source          destination
```

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
```

Chain POSTROUTING (policy ACCEPT)

```
target  prot opt source          destination
```

Table: filter

Chain INPUT (policy ACCEPT)

target	prot	opt	source	destination
RH-Firewall-1-INPUT	all	--	0.0.0.0/0	0.0.0.0/0

Chain FORWARD (policy ACCEPT)

target	prot	opt	source	destination
RH-Firewall-1-INPUT	all	--	0.0.0.0/0	0.0.0.0/0

Chain OUTPUT (policy ACCEPT)

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

Chain RH-Firewall-1-INPUT (2 references)

target	prot	opt	source	destination	
ACCEPT	all	--	0.0.0.0/0	0.0.0.0/0	
ACCEPT	icmp	--	0.0.0.0/0	0.0.0.0/0	icmp type 255
ACCEPT	esp	--	0.0.0.0/0	0.0.0.0/0	
ACCEPT	ah	--	0.0.0.0/0	0.0.0.0/0	
ACCEPT	udp	--	0.0.0.0/0	224.0.0.251	udp dpt:5353
ACCEPT	udp	--	0.0.0.0/0	0.0.0.0/0	udp dpt:631
ACCEPT	all	--	0.0.0.0/0	0.0.0.0/0	state RELATED,ESTABLISHED
ACCEPT	tcp	--	0.0.0.0/0	0.0.0.0/0	state NEW tcp dpt:80
ACCEPT	tcp	--	0.0.0.0/0	0.0.0.0/0	state NEW tcp dpt:443
ACCEPT	tcp	--	0.0.0.0/0	0.0.0.0/0	state NEW tcp dpt:22
ACCEPT	tcp	--	0.0.0.0/0	0.0.0.0/0	state NEW tcp dpt:25
ACCEPT	tcp	--	0.0.0.0/0	0.0.0.0/0	state NEW tcp dpt:139
ACCEPT	tcp	--	0.0.0.0/0	0.0.0.0/0	state NEW tcp dpt:445
REJECT	all	--	0.0.0.0/0	0.0.0.0/0	reject-with icmp-host-prohibited

Both the host and guest OS are CentOS 4.

Any ideas other than turn off the firewall?

jch

Subject: Re: DNS not working in VE
Posted by [dim](#) on Thu, 13 Jul 2006 08:13:12 GMT
[View Forum Message](#) <> [Reply to Message](#)

DNS uses udp, port=53. I don't see any related rule in your output. So, you need to add something like this to your rules:

```
iptables -A RH-Firewall-1-INPUT -p udp --source-port 53 -j ACCEPT
iptables -A RH-Firewall-1-INPUT -p udp --destination-port 53 -j ACCEPT
```

Subject: Re: DNS not working in VE

Posted by [jchamilton](#) on Thu, 13 Jul 2006 13:35:35 GMT

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

Didn't work - here's what I did:

Added the rules by copy/paste

```
# service iptables save
```

```
# service vz stop
```

```
# service iptables restart
```

```
# service vz start
```

```
# vzctl start 101
```

```
# vzctl enter 101
```

Then from the VE:

```
entered into VPS 101
```

```
-bash-3.00# ping www.google.com
```

```
ping: unknown host www.google.com
```

With strace:

```
execve("/bin/ping", ["ping", "www.google.com"], [/* 16 vars */]) = 0
```

```
uname({sys="Linux", node="vps101.corp.xxxxx.com", ...}) = 0
```

```
brk(0) = 0x8001c000
```

```
access("/etc/ld.so.preload", R_OK) = -1 ENOENT (No such file or directory)
```

```
open("/etc/ld.so.cache", O_RDONLY) = 3
```

```
fstat64(3, {st_mode=S_IFREG|0644, st_size=18909, ...}) = 0
```

```
old_mmap(NULL, 18909, PROT_READ, MAP_PRIVATE, 3, 0) = 0x40017000
```

```
close(3) = 0
```

```
open("/lib/libresolv.so.2", O_RDONLY) = 3
```

```
read(3, "\177ELF\1\1\0\0\0\0\0\0\0\0\3\0\3\0\1\0\0\0\360#\0"..., 512) = 512
```

```
fstat64(3, {st_mode=S_IFREG|0755, st_size=77740, ...}) = 0
```

```
old_mmap(NULL, 75944, PROT_READ|PROT_EXEC, MAP_PRIVATE|MAP_DENYWRITE, 3, 0)
```

```
= 0x4001c000
```

```
old_mmap(0x4002b000, 8192, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_FIXED|MAP_DENYWRITE, 3, 0xf000) = 0x4002b000
```

```
old_mmap(0x4002d000, 6312, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_FIXED|MAP_ANONYMOUS, -1, 0) = 0x4002d000
```

```
close(3) = 0
```

```
open("/lib/tls/libc.so.6", O_RDONLY) = 3
```

```
read(3, "\177ELF\1\1\0\0\0\0\0\0\0\0\3\0\3\0\1\0\0\0\200\1\000"..., 512) = 512
```

```
fstat64(3, {st_mode=S_IFREG|0755, st_size=1451450, ...}) = 0
```

```
old_mmap(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1, 0) = 0x4002f000
```

```
old_mmap(NULL, 1219772, PROT_READ|PROT_EXEC, MAP_PRIVATE|MAP_DENYWRITE, 3, 0) = 0x40030000
```

```
old_mmap(0x40154000, 16384, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_FIXED|MAP_DENYWRITE, 3, 0x124000) = 0x40154000
```

```

old_mmap(0x40158000, 7356, PROT_READ|PROT_WRITE,
MAP_PRIVATE|MAP_FIXED|MAP_ANONYMOUS, -1, 0) = 0x40158000
close(3) = 0
old_mmap(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1,
0) = 0x4015a000
mprotect(0x40154000, 4096, PROT_READ) = 0
mprotect(0x40015000, 4096, PROT_READ) = 0
set_thread_area({entry_number:-1 -> 6, base_addr:0x4015a6c0, limit:1048575, seg_32bit:1,
contents:0, read_exec_only:0, limit_in_pages:1, seg_not_present:0, useable:1}) = 0
munmap(0x40017000, 18909) = 0
socket(PF_INET, SOCK_RAW, IPPROTO_ICMP) = 3
getuid32() = 0
setuid32(0) = 0
brk(0) = 0x8001c000
brk(0x8003d000) = 0x8003d000
gettimeofday({1152796486, 371270}, NULL) = 0
getpid() = 7922
open("/etc/resolv.conf", O_RDONLY) = 4
fstat64(4, {st_mode=S_IFREG|0644, st_size=23, ...}) = 0
mmap2(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1, 0) =
0x40017000
read(4, "nameserver 10.79.32.60\n", 4096) = 23
read(4, "", 4096) = 0
close(4) = 0
munmap(0x40017000, 4096) = 0
uname({sys="Linux", node="vps101.corp.xxxxx.com", ...}) = 0
socket(PF_FILE, SOCK_STREAM, 0) = 4
fcntl64(4, F_GETFL) = 0x2 (flags O_RDWR)
fcntl64(4, F_SETFL, O_RDWR|O_NONBLOCK) = 0
connect(4, {sa_family=AF_FILE, path="/var/run/nscd/socket"}, 110) = -1 ENOENT (No such file or
directory)
close(4) = 0
socket(PF_FILE, SOCK_STREAM, 0) = 4
fcntl64(4, F_GETFL) = 0x2 (flags O_RDWR)
fcntl64(4, F_SETFL, O_RDWR|O_NONBLOCK) = 0
connect(4, {sa_family=AF_FILE, path="/var/run/nscd/socket"}, 110) = -1 ENOENT (No such file or
directory)
close(4) = 0
open("/etc/nsswitch.conf", O_RDONLY) = 4
fstat64(4, {st_mode=S_IFREG|0644, st_size=1623, ...}) = 0
mmap2(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1, 0) =
0x40017000
read(4, "#\n# /etc/nsswitch.conf\n#\n# An ex"..., 4096) = 1623
read(4, "", 4096) = 0
close(4) = 0
munmap(0x40017000, 4096) = 0
open("/etc/ld.so.cache", O_RDONLY) = 4
fstat64(4, {st_mode=S_IFREG|0644, st_size=18909, ...}) = 0

```

```
old_mmap(NULL, 18909, PROT_READ, MAP_PRIVATE, 4, 0) = 0x40017000
close(4) = 0
open("/lib/libnss_files.so.2", O_RDONLY) = 4
read(4, "\177ELF\1\1\1\0\0\0\0\0\0\0\0\3\0\3\0\1\0\0\0\20\33\0"..., 512) = 512
fstat64(4, {st_mode=S_IFREG|0755, st_size=45889, ...}) = 0
old_mmap(NULL, 41608, PROT_READ|PROT_EXEC, MAP_PRIVATE|MAP_DENYWRITE, 4, 0) = 0x4015b000
old_mmap(0x40164000, 8192, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_FIXED|MAP_DENYWRITE, 4, 0x8000) = 0x40164000
close(4) = 0
munmap(0x40017000, 18909) = 0
open("/etc/host.conf", O_RDONLY) = 4
fstat64(4, {st_mode=S_IFREG|0644, st_size=17, ...}) = 0
mmap2(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1, 0) = 0x40017000
read(4, "order hosts,bind\n", 4096) = 17
read(4, "", 4096) = 0
close(4) = 0
munmap(0x40017000, 4096) = 0
open("/etc/hosts", O_RDONLY) = 4
fcntl64(4, F_GETFD) = 0
fcntl64(4, F_SETFD, FD_CLOEXEC) = 0
fstat64(4, {st_mode=S_IFREG|0644, st_size=73, ...}) = 0
mmap2(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1, 0) = 0x40017000
read(4, "127.0.0.1 vps101.corp.xxxxx.co"..., 4096) = 73
read(4, "", 4096) = 0
close(4) = 0
munmap(0x40017000, 4096) = 0
open("/etc/ld.so.cache", O_RDONLY) = 4
fstat64(4, {st_mode=S_IFREG|0644, st_size=18909, ...}) = 0
old_mmap(NULL, 18909, PROT_READ, MAP_PRIVATE, 4, 0) = 0x40017000
close(4) = 0
open("/lib/libnss_dns.so.2", O_RDONLY) = 4
read(4, "\177ELF\1\1\1\0\0\0\0\0\0\0\0\3\0\3\0\1\0\0\0\300\r\0"..., 512) = 512
fstat64(4, {st_mode=S_IFREG|0755, st_size=21280, ...}) = 0
old_mmap(NULL, 20612, PROT_READ|PROT_EXEC, MAP_PRIVATE|MAP_DENYWRITE, 4, 0) = 0x40166000
old_mmap(0x4016a000, 8192, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_FIXED|MAP_DENYWRITE, 4, 0x3000) = 0x4016a000
close(4) = 0
munmap(0x40017000, 18909) = 0
socket(PF_INET, SOCK_DGRAM, IPPROTO_IP) = 4
connect(4, {sa_family=AF_INET, sin_port=htons(53), sin_addr=inet_addr("10.79.32.60")}, 28) = 0
fcntl64(4, F_GETFL) = 0x2 (flags O_RDWR)
fcntl64(4, F_SETFL, O_RDWR|O_NONBLOCK) = 0
gettimeofday({1152796486, 377655}, NULL) = 0
poll([{fd=4, events=POLLOUT, revents=POLLOUT}], 1, 0) = 1
```

```
send(4, "\363\362\1\0\0\1\0\0\0\0\0\0\0\3www\6google\3com\0\0\1\0"... , 32, 0) = 32
poll([fd=4, events=POLLIN, revents=POLLERR]), 1, 5000) = 1
close(4) = 0
socket(PF_INET, SOCK_DGRAM, IPPROTO_IP) = 4
connect(4, {sa_family=AF_INET, sin_port=htons(53), sin_addr=inet_addr("10.79.32.60")}, 28) = 0
fcntl64(4, F_GETFL) = 0x2 (flags O_RDWR)
fcntl64(4, F_SETFL, O_RDWR|O_NONBLOCK) = 0
gettimeofday({1152796486, 378628}, NULL) = 0
poll([fd=4, events=POLLOUT, revents=POLLOUT]), 1, 0) = 1
send(4, "\363\362\1\0\0\1\0\0\0\0\0\0\0\3www\6google\3com\0\0\1\0"... , 32, 0) = 32
poll([fd=4, events=POLLIN, revents=POLLERR]), 1, 5000) = 1
close(4) = 0
socket(PF_INET, SOCK_DGRAM, IPPROTO_IP) = 4
connect(4, {sa_family=AF_INET, sin_port=htons(53), sin_addr=inet_addr("10.79.32.60")}, 28) = 0
fcntl64(4, F_GETFL) = 0x2 (flags O_RDWR)
fcntl64(4, F_SETFL, O_RDWR|O_NONBLOCK) = 0
gettimeofday({1152796486, 379496}, NULL) = 0
poll([fd=4, events=POLLOUT, revents=POLLOUT]), 1, 0) = 1
send(4, "L_\1\0\0\1\0\0\0\0\0\0\0\3www\6google\3com\4corp"... , 48, 0) = 48
poll([fd=4, events=POLLIN, revents=POLLERR]), 1, 5000) = 1
close(4) = 0
socket(PF_INET, SOCK_DGRAM, IPPROTO_IP) = 4
connect(4, {sa_family=AF_INET, sin_port=htons(53), sin_addr=inet_addr("10.79.32.60")}, 28) = 0
fcntl64(4, F_GETFL) = 0x2 (flags O_RDWR)
fcntl64(4, F_SETFL, O_RDWR|O_NONBLOCK) = 0
gettimeofday({1152796486, 380237}, NULL) = 0
poll([fd=4, events=POLLOUT, revents=POLLOUT]), 1, 0) = 1
send(4, "L_\1\0\0\1\0\0\0\0\0\0\0\3www\6google\3com\4corp"... , 48, 0) = 48
poll([fd=4, events=POLLIN, revents=POLLERR]), 1, 5000) = 1
close(4) = 0
write(2, "ping: unknown host www.google.co"... , 34ping: unknown host www.google.com
) = 34
exit_group(2) = ?
Process 7922 detached
```

I also uninstalled the only package I've installed since I installed OpenVZ with the same results...

jch

Subject: Re: DNS not working in VE
Posted by [jchamilton](#) on Thu, 13 Jul 2006 14:19:31 GMT
[View Forum Message](#) <> [Reply to Message](#)

Fixed it - it was indeed the firewall rule but I recalled that the rules are taken in order and when I added the rules they got appended thusly:

Chain RH-Firewall-1-INPUT (2 references)

```
target  prot opt source          destination
ACCEPT  all  --  0.0.0.0/0      0.0.0.0/0
ACCEPT  icmp --  0.0.0.0/0      0.0.0.0/0      icmp type 255
ACCEPT  esp  --  0.0.0.0/0      0.0.0.0/0
ACCEPT  ah   --  0.0.0.0/0      0.0.0.0/0
ACCEPT  udp  --  0.0.0.0/0      224.0.0.251    udp dpt:5353
ACCEPT  udp  --  0.0.0.0/0      0.0.0.0/0      udp dpt:631
ACCEPT  all  --  0.0.0.0/0      0.0.0.0/0      state RELATED,ESTABLISHED
ACCEPT  tcp  --  0.0.0.0/0      0.0.0.0/0      state NEW tcp dpt:80
ACCEPT  tcp  --  0.0.0.0/0      0.0.0.0/0      state NEW tcp dpt:443
ACCEPT  tcp  --  0.0.0.0/0      0.0.0.0/0      state NEW tcp dpt:22
ACCEPT  tcp  --  0.0.0.0/0      0.0.0.0/0      state NEW tcp dpt:25
ACCEPT  tcp  --  0.0.0.0/0      0.0.0.0/0      state NEW tcp dpt:139
ACCEPT  tcp  --  0.0.0.0/0      0.0.0.0/0      state NEW tcp dpt:445
REJECT  all  --  0.0.0.0/0      0.0.0.0/0      reject-with icmp-host-prohibited
ACCEPT  udp  --  0.0.0.0/0      0.0.0.0/0      udp spt:53
ACCEPT  udp  --  0.0.0.0/0      0.0.0.0/0      udp dpt:53
```

So I edited /etc/sysconfig/iptables and put the accept targets for udp before the 'REJECT all' line:

Chain RH-Firewall-1-INPUT (2 references)

```
target  prot opt source          destination
ACCEPT  all  --  0.0.0.0/0      0.0.0.0/0
ACCEPT  icmp --  0.0.0.0/0      0.0.0.0/0      icmp type 255
ACCEPT  esp  --  0.0.0.0/0      0.0.0.0/0
ACCEPT  ah   --  0.0.0.0/0      0.0.0.0/0
ACCEPT  udp  --  0.0.0.0/0      224.0.0.251    udp dpt:5353
ACCEPT  udp  --  0.0.0.0/0      0.0.0.0/0      udp dpt:631
ACCEPT  all  --  0.0.0.0/0      0.0.0.0/0      state RELATED,ESTABLISHED
ACCEPT  tcp  --  0.0.0.0/0      0.0.0.0/0      state NEW tcp dpt:80
ACCEPT  tcp  --  0.0.0.0/0      0.0.0.0/0      state NEW tcp dpt:443
ACCEPT  tcp  --  0.0.0.0/0      0.0.0.0/0      state NEW tcp dpt:22
ACCEPT  tcp  --  0.0.0.0/0      0.0.0.0/0      state NEW tcp dpt:25
ACCEPT  tcp  --  0.0.0.0/0      0.0.0.0/0      state NEW tcp dpt:139
ACCEPT  tcp  --  0.0.0.0/0      0.0.0.0/0      state NEW tcp dpt:445
ACCEPT  udp  --  0.0.0.0/0      0.0.0.0/0      udp spt:53
ACCEPT  udp  --  0.0.0.0/0      0.0.0.0/0      udp dpt:53
REJECT  all  --  0.0.0.0/0      0.0.0.0/0      reject-with icmp-host-prohibited
```

Restarted iptables, and it worked!

Thanks for the help dim!

jch
