Subject: OpenVZ 6.3 - No VPS have Connectivity until "service network restart" Posted by MikeDVB on Sun, 18 Nov 2012 00:53:38 GMT View Forum Message <> Reply to Message

I just set up a brand new OpenVZ 6.3 server, disabled SELINUX, and migrated a single VPS over to the node using vzmigrate. That went fine, but once the VPS was migrated over the VPS has no connectivity to the outside world.

I can ping the VPS' IP from inside the node, and from another server on the same switch, but I cannot reach it from outside of our network nor can the VPS reach anything but the node itself. It can't ping google, or even a direct IP.

If I run 'service network restart' on the node, the VPS will suddenly have connectivity. This is without making *any* changes to *anything* after the reboot and before doing 'service network restart'. If I do 'service network restart' and the VPS has connectivity - on reboot it no longer has connectivity until I restart the network manually again.

service network restart

Shutting down interface eth0:[OK]Shutting down interface eth1:[OK]Shutting down interface venet0:Shutting down interface venet0:

	[OK]
Shutting down loopback interface:	[OK]
Bringing up loopback interface:	[OK]
Bringing up interface eth0:	[OK]
Bringing up interface eth1:	[OK]
Bringing up interface venet0: Bringing	ng up interface venet0:
Configuring interface venet0:	
net.ipv4.conf.venet0.send_redirects = 0	
Configuring ipv6 venet0:	

[OK]

I do see from the boot.log that it is doing the network start-up using the same /etc/init.d/network script:

cat boot.log Welcome to CentOS Starting udev: [OK] Setting hostname redacted.redacted.com: [OK] Checking filesystems /dev/sdb2: clean, 63294/182255616 files, 18157450/729021952 blocks /dev/sda1: clean, 50/131072 files, 85436/523264 blocks [OK] Remounting root filesystem in read-write mode: [OK] Mounting local filesystems: [OK] Enabling /etc/fstab swaps: [OK] Entering non-interactive startup Calling the system activity data collector (sadc): [OK] ip6tables: Applying firewall rules: iptables: Applying firewall rules: [OK]

iptables: Loading additional modules: ipt_REJECT ipt_tos ipt_TOS ipt_LOG ip_conntrack ipt_limit ipt multiport iptable filter iptable mangle ipt TCPMSS ipt tcpmss ipt ttl ipt length ipt state ip[OK]t ip_nat_ftp ipt_owner ipt_REDIRECT Bringing up loopback interface: [OK] Bringing up interface eth0: [OK] Bringing up interface eth1: [OK] Bringing up interface venet0: Bringing up interface venet0: Configuring interface venet0: net.ipv4.conf.venet0.send redirects = 0 Configuring ipv6 venet0: [OK] Starting auditd: [OK] [OK] Starting portreserve: Starting system logger: [OK] Mounting other filesystems: [OK] Retrigger failed udev events [OK] Starting sshd: [OK] Starting postfix: [OK] Starting lighttpd: [OK] Starting... Starting... Starting crond: [OK] Checking vzevent kernel module ... [OK] [OK] Starting vzeventd: Starting OpenVZ: [OK] Applying OOM adjustments: [OK] Starting CT 258: [OK] The VPS is indeed running after the boot: # vzlist -a CTID NPROC STATUS IP ADDR HOSTNAME 258 36 running re.da.ct.ed redacted.redacted.com ipv4.ip_forward is enabled: # sysctl -e net.ipv4.ip_forward net.ipv4.ip forward = 1I haven't touched the /etc/vz/vz.conf: # cat /etc/vz/vz.conf

cat /etc/vz/vz.conf ## Global parameters VIRTUOZZO=yes LOCKDIR=/vz/lock DUMPDIR=/vz/dump VE0CPUUNITS=1000

Logging parameters LOGGING=yes LOGFILE=/var/log/vzctl.log LOG_LEVEL=0 VERBOSE=0

Disk quota parameters DISK_QUOTA=yes VZFASTBOOT=no

Disable module loading. If set, vz initscript does not load any modules. #MODULES_DISABLED=yes

The name of the device whose IP address will be used as source IP for CT.
By default automatically assigned.
#VE_ROUTE_SRC_DEV="eth0"

Controls which interfaces to send ARP requests and modify APR tables on. NEIGHBOUR_DEVS=all

Fail if there is another machine in the network with the same IP ERROR_ON_ARPFAIL="no"

Template parameters TEMPLATE=/vz/template

Defaults for containers VE_ROOT=/vz/root/\$VEID VE_PRIVATE=/vz/private/\$VEID CONFIGFILE="basic" DEF_OSTEMPLATE="centos-5"

Load vzwdog module VZWDOG="no"

IPv4 iptables kernel modules to be enabled in CTs by default IPTABLES="ipt_REJECT ipt_tos ipt_TOS ipt_LOG ip_conntrack ipt_limit ipt_multiport iptable_filter iptable_mangle ipt_TCPMSS ipt_tcpmss ipt_ttl ipt_length ipt_state iptable_nat ip_nat_ftp ipt_owner ipt_REDIRECT" ## IPv4 iptables kernel modules to be loaded by init.d/vz script IPTABLES_MODULES="\$IPTABLES"

Enable IPv6 IPV6="yes"

IPv6 ip6tables kernel modules IP6TABLES="ip6_tables ip6table_filter ip6table_mangle ip6t_REJECT"

I flushed iptables rules on the main node, which didn't help/solve the issue.

I'm really at a loss as to what would cause this. I thought it may be the arp table, but if that were

the case - a reboot (takes less than 2 minutes) should not cause failures - especially since the ARP entries don't time out that fast on our network

Traceroutes from external sources and other servers show that the traffic is hitting the node before timing out:

traceroute re.da.ct.ed

traceroute to re.da.ct.ed (re.da.ct.ed), 30 hops max, 40 byte packets

1 the-node.redacted.com (the-node.primary-ip.com) 0.150 ms 0.133 ms 0.121 ms

2 *** 3 ***

etc...

The VPS can ping the node's IP, and the node can ping the VPS' IP before and after 'service network restart'.