


BalanceNG[®] V3

A Modern Software Load Balancer for Linux and Solaris

Proof of Concept

Connecting to local TCP/IP stack using BNG “tap” interfaces



Status: **DRAFT**
Date: Feb 11, 2012
Author: Thomas Obermair



Inlab Software GmbH
Josef-Würth-Str. 3
82031 Grünwald
Germany

Tel.: +49 89 64911420
Fax: +49 89 6411160
Email: office@inlab.de
Home: <http://www.inlab.de>

Table of Contents

1 Rationale.....	4
2 rp_filter.....	4
3 loopback alias.....	4
4 ARP flux prevention.....	4
5 bng.conf.....	4
6 POC.....	5

Legal Notices

© Copyright 2005-2011, 2012 by Inlab Software GmbH, Josef-Wuerth-Str. 3, Gruenwald, Germany. All Rights Reserved / Alle Rechte vorbehalten.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Inlab Software GmbH.

BalanceNG is a trademark of Inlab Software GmbH. Gentoo is a trademark by Gentoo Technologies, Inc. Debian is a registered trademark of Software In The Public Interest, Inc. FreeBSD is a registered trademark of Walnut Creek CDROM, Inc. Linux is a registered trademark of Linus Torvalds. All other trademarks and registered trademarks mentioned in this document are properties by their respective holders.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

1 Rationale

This proof of concept shows how to connect to the local host OS TCP/IP stack using BNG “tap” interfaces and DSR simultaneously. This technique is useful to connect to local running co-processes which may implement SSL offloading, HTTP header inspection, SIP multiplexing/demultiplexing and many other useful functionalities.

2 rp_filter

The Linux “rp_filter” Kernel parameters needs to be set as follows:

```
# for i in /proc/sys/net/ipv4/conf/*/rp_filter ; do echo 0 > $i ; done
```

3 loopback alias

The loopback alias needs to be established as follows (using the virtual server 1 IP address):

```
# ifconfig lo:0 172.17.2.90 netmask 255.255.255.255 -arp up
```

4 ARP flux prevention

The following two Kernel parameters need to be set as follows (as usual and needed for DSR):

```
# echo 1 > /proc/sys/net/ipv4/conf/all/arp_ignore
# echo 2 > /proc/sys/net/ipv4/conf/all/arp_announce
```

5 bng.conf

```
// configuration taken Fri Feb 10 20:02:35 2012
// BalanceNG 3.373 (created 2012/02/08)
// INLAB 39ed5ff716d8a118a7c61f777cbb0446
modules vrrp,arp,ping,hc,master,slb,tnat,nat,rt
set localdsr 1
interface 1 {
    name eth0
    access raw
}
interface 2 {
    name bng0
    access tap
    init "ip addr add 10.20.20.20/24 dev bng0; ip link set bng0 up"
}
register interfaces 1,2
enable interfaces 1,2
vrrp {
    vrid 79
    priority 200
    network 1
}
network 1 {
    addr 172.17.2.0
    mask 255.255.255.0
    real 172.17.2.55
    virt 172.17.2.56
    interface 1
}
network 2 {
    addr 10.20.20.0
    mask 255.255.255.0
    real 10.20.20.1
    virt 10.20.20.2
    interface 2
}
register networks 1,2
enable networks 1,2
server 1 {
    ipaddr 172.17.2.90
    port 22
    protocol tcp
    target 1
}
```

```
register server 1
enable server 1
target 1 {
    ipaddr 10.20.20.20
    port 22
    protocol tcp
    ping 2,10
    dsr enable
}
register target 1
enable target 1
// end of configuration
```

6 POC

Connecting with telnet from a third test client succeeds immediately (thus connecting to the local sshd):

```
t@src:~ $ telnet 172.17.2.90 22
Trying 172.17.2.90...
Connected to 172.17.2.90.
Escape character is '^]'.
SSH-2.0-OpenSSH_5.8p1 Debian-7ubuntu1

Protocol mismatch.
Connection to 172.17.2.90 closed by foreign host.
t@src:~ $
```

The expected session table entry has been created:

```
# bng control
BalanceNG: connected to PID 25512
bng# sh sessions
0 sessions
bng# sh sessions
1 session
  srv tgt  age timeout ftimeout SYNC session-id
  ---
  1 1 62 600 0 172.17.2.92:0
bng#
```