

By Falko Timme

Published: 2009-01-04 17:28

Setting Up A High-Availability Load Balancer (With Failover and Session Support) With HAProxy/Wackamole/Spread On Debian Etch

Version 1.0

Author: Falko Timme <ft [at] falkotimme [dot] com>

Last edited 12/22/2008

This article explains how to set up a two-node load balancer in an active/passive configuration with [HAProxy](#), [Wackamole](#), and [Spread](#) on Debian Etch. The load balancer sits between the user and two (or more) backend Apache web servers that hold the same content. Not only does the load balancer distribute the requests to the two backend Apache servers, it also checks the health of the backend servers. If one of them is down, all requests will automatically be redirected to the remaining backend server. In addition to that, the two load balancer nodes monitor each other using Wackamole and Spread, and if the master fails, the slave becomes the master, which means the users will not notice any disruption of the service. HAProxy is session-aware, which means you can use it with any web application that makes use of sessions (such as forums, shopping carts, etc.).

From the HAProxy web site: *"HAProxy is a free, very fast and reliable solution offering high availability, load balancing, and proxying for TCP and HTTP-based applications. It is particularly suited for web sites crawling under very high loads while needing persistence or Layer7 processing. Supporting tens of thousands of connections is clearly realistic with todays hardware. Its mode of operation makes its integration into existing architectures very easy and riskless, while still offering the possibility not to expose fragile web servers to the Net."*

I do not issue any guarantee that this will work for you!

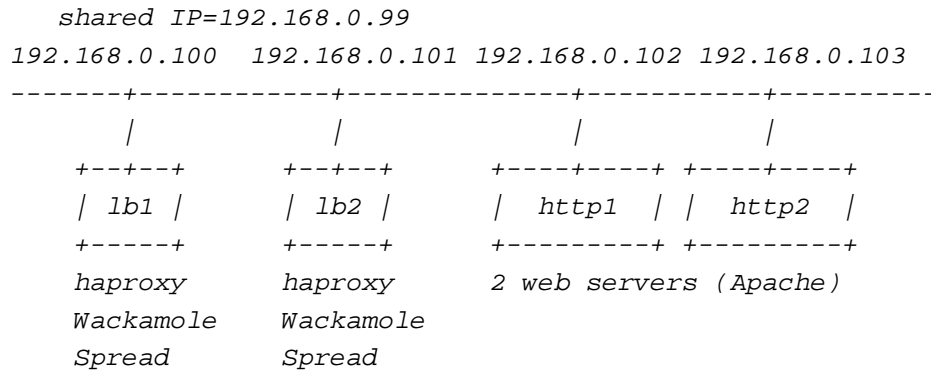
1 Preliminary Note

In this tutorial I will use the following hosts:

- Load Balancer 1: `lb1.example.com`, IP address: `192.168.0.100`
- Load Balancer 2: `lb2.example.com`, IP address: `192.168.0.101`
- Web Server 1: `http1.example.com`, IP address: `192.168.0.102`

- Web Server 2: `http2.example.com`, IP address: `192.168.0.103`
- We also need a virtual IP address that floats between `lb1` and `lb2`: `192.168.0.99`

Here's a little diagram that shows our setup:



The shared (virtual) IP address is no problem as long as you're in your own LAN where you can assign IP addresses as you like. However, if you want to use this setup with public IP addresses, you need to find a hoster where you can rent two servers (the load balancer nodes) in the same subnet; you can then use a free IP address in this subnet for the virtual IP address.

`http1` and `http2` are standard Debian Etch Apache setups with the document root `/var/www` (the configuration of this default vhost is stored in `/etc/apache2/sites-available/default`). If your document root differs, you might have to adjust this guide a bit.

To make HAProxy session-aware, I'm assuming that the web application that is installed on `http1` and `http2` uses the session id `JSESSIONID`.

2 Preparing The Backend Web Servers

We will configure HAProxy as a transparent proxy, i.e., it will pass on the original user's IP address in a field called `X-Forwarded-For` to the backend web servers. Of course, the backend web servers should log the original user's IP address in their access logs instead of the IP addresses of our load balancers. Therefore we must modify the `LogFormat` line in `/etc/apache2/apache2.conf` and replace `%h` with `%{X-Forwarded-For}i`:

[http1/http2:](#)

```
vi /etc/apache2/apache2.conf
```

```
[...]  
#LogFormat "%h %l %u %t \"%r\" %>s %b \"%{Referer}i\" \"%{User-Agent}i\"" combined  
LogFormat "%{X-Forwarded-For}i %l %u %t \"%r\" %>s %b \"%{Referer}i\" \"%{User-Agent}i\"" combined  
[...]
```

Also, we will configure HAProxy to check the backend servers' health by continuously requesting the file `check.txt` (translates to `/var/www/check.txt` if `/var/www` is your document root) from the backend servers. Of course, these requests would totally bloat the access logs and mess up your page view statistics (if you use a tool like Webalizer or AWstats that generates statistics based on the access logs).

Therefore we open our vhost configuration (in this example it's in `/etc/apache2/sites-available/default`) and put these two lines into it (comment out all other `CustomLog` directives in your vhost configuration):

```
vi /etc/apache2/sites-available/default
```

```
[...]  
SetEnvIf Request_URI "^/check\.txt$" dontlog  
CustomLog /var/log/apache2/access.log combined env=!dontlog  
[...]
```

This configuration prevents that requests to `check.txt` get logged in Apache's access log.

Afterwards we restart Apache:

```
/etc/init.d/apache2 restart
```

... and create the file `check.txt` (this can be an empty file):

```
touch /var/www/check.txt
```

We are finished already with the backend servers; the rest of the configuration happens on the two load balancer nodes.

3 Installing HAProxy

Unfortunately HAProxy is available as a Debian package for Debian Lenny (testing) and Sid (unstable), but not for Etch. Therefore we will install the HAProxy package from Lenny. To do this, open `/etc/apt/sources.list` and add the line `deb http://ftp2.de.debian.org/debian/ lenny main`; your `/etc/apt/sources.list` could then look like this:

[lb1/lb2](#):

```
vi /etc/apt/sources.list
```

```
deb http://ftp2.de.debian.org/debian/ etch main
deb-src http://ftp2.de.debian.org/debian/ etch main

deb http://ftp2.de.debian.org/debian/ lenny main

deb http://security.debian.org/ etch/updates main contrib
deb-src http://security.debian.org/ etch/updates main contrib
```

Of course (in order not to mess up our system), we want to install packages from Lenny only if there's no appropriate package from Etch - if there are packages from Etch and Lenny, we want to install the one from Etch. To do this, we give packages from Etch a higher priority in `/etc/apt/preferences`:

```
vi /etc/apt/preferences
```

```
Package: *  
Pin: release a=etch  
Pin-Priority: 700  
  
Package: *  
Pin: release a=lenny  
Pin-Priority: 650
```

(The terms *etch* and *lenny* refer to the appropriate terms in `/etc/apt/sources.list`; if you're using *stable* and *testing* there, you must use *stable* and *testing* instead of *etch* and *lenny* in `/etc/apt/preferences` as well.)

Afterwards, we update our packages database:

```
apt-get update
```

(If you get an error like this one:

```
E: Dynamic MMap ran out of room
```

then open `/etc/apt/apt.conf...`

```
vi /etc/apt/apt.conf
```

... and add a line for `APT::Cache-Limit` with a very high value, e.g. like this:

```
APT::Cache-Limit "100000000";
```

Then run

```
apt-get update
```

again.)

Upgrade the installed packages:

```
apt-get upgrade
```

... and install HAProxy:

```
apt-get install haproxy
```

4 Configuring The Load Balancers

The HAProxy configuration is stored in `/etc/haproxy/haproxy.cfg` and is pretty straight-forward. I won't explain all the directives here; to learn more about all options, please read <http://haproxy.1wt.eu/download/1.3/doc/haproxy-en.txt> and <http://haproxy.1wt.eu/download/1.2/doc/architecture.txt>.

We back up the original `/etc/haproxy/haproxy.cfg` and create a new one like this:

lb1/lb2:

```
cp /etc/haproxy/haproxy.cfg /etc/haproxy/haproxy.cfg_orig  
  
cat /dev/null > /etc/haproxy/haproxy.cfg  
  
vi /etc/haproxy/haproxy.cfg
```

```
global
  log 127.0.0.1 local0
  log 127.0.0.1 local1 notice
  #log loghost local0 info
  maxconn 4096
  #debug
  #quiet
  user haproxy
  group haproxy

defaults
  log global
  mode http
  option httplog
  option dontlognull
  retries 3
  redispatch
  maxconn 2000
  contimeout 5000
  clitimeout 50000
  srvtimeout 50000

listen webfarm 192.168.0.99:80
  mode http
  stats enable
  stats auth someuser:somepassword
  balance roundrobin
  cookie JSESSIONID prefix
  option httpclose
  option forwardfor
  option httpchk HEAD /check.txt HTTP/1.0
  server webA 192.168.0.102:80 cookie A check
  server webB 192.168.0.103:80 cookie B check
```

Afterwards, we set *ENABLED* to 1 in */etc/default/haproxy*:

```
vi /etc/default/haproxy
```

```
# Set ENABLED to 1 if you want the init script to start haproxy.  
ENABLED=1  
# Add extra flags here.  
#EXTRAOPTS="-de -m 16"
```

5 Setting Up Wackamole/Spread

We've just configured HAProxy to listen on the virtual IP address *192.168.0.99*, but someone has to tell *lb1* and *lb2* that they should listen on that IP address. This is done by Wackamole and Spread which we install like this:

[lb1/lb2:](#)

```
apt-get install wackamole
```

To allow HAProxy to bind to the shared IP address, we add the following line to */etc/sysctl.conf*:

```
vi /etc/sysctl.conf
```

```
[...]  
net.ipv4.ip_nonlocal_bind=1
```


... and run:

```
sysctl -p
```

Next we modify `/etc/default/spread` and set `ENABLED` to 1:

```
vi /etc/default/spread
```

```
# Change to enable spread
ENABLED=1

# Options, see spread.1 for list
OPTIONS=""
```

The Spread configuration is located in `/etc/spread/spread.conf`. We create a backup of the original file and edit `spread.conf` as follows:

```
cp /etc/spread/spread.conf /etc/spread/spread.conf_orig
cat /dev/null > /etc/spread/spread.conf
vi /etc/spread/spread.conf
```

Spread can send broadcast or multicast messages, therefore you have two choices for configuring Spread.

[Option 1 \(Broadcast Messages\):](#)

```
Spread_Segment 192.168.0.255 {
```

```
lb1.example.com 192.168.0.100
lb2.example.com 192.168.0.101
}

EventLogFile = /var/log/spread.log

EventTimeStamp
```

Option 2 (Multicast Messages):

```
Spread_Segment 225.0.1.1 {
  lb1.example.com 192.168.0.100
  lb2.example.com 192.168.0.101
}

EventLogFile = /var/log/spread.log

EventTimeStamp
```

Spread will feel free to use broadcast messages within a sub-network if you use broadcast messages. If IP-multicast is supported by the operating system, then the messages will only be received by those machines who are in the group and not by all others in the same sub-network as happens with broadcast addresses

Now we can start Spread:

```
/etc/init.d/spread start
```

On to the Wackamole configuration. Open `/etc/default/wackamole` and set `ENABLED` to 1:

```
vi /etc/default/wackamole
```

```
# Change to enable wackamole
ENABLED=1

# Options
OPTIONS=""
```

Then configure Wackamole as follows:

```
cp /etc/wackamole.conf /etc/wackamole.conf_orig

cat /dev/null > /etc/wackamole.conf

vi /etc/wackamole.conf
```

```
Spread = 4803
SpreadRetryInterval = 5s
Group = wack1
Control = /var/run/wackamole/wackamole.it

Prefer None

VirtualInterfaces {
    eth0:192.168.0.99/24
}

Arp-Cache = 60s
```

```
Notify {
    eth0:192.168.0.1/32
    eth0:192.0.0.0/24
    arp-cache
}
balance {
    AcquisitionsPerRound = all
    interval = 4s
}
mature = 5s
```

The *VirtualInterfaces* stanza is the most important part - it contains our network interface (*eth0* in this example) together with our virtual IP address (*192.168.0.99*).

The *Notify* stanza contains hosts (e.g. your router) or subnets to notify when the virtual IP switches. It is not necessary, so you can leave out that stanza, if you like.

(You can learn more about the Wackamole configuration by taking a look at

```
man 5 wackamole.conf
```

)

Finally we start Wackamole on both load balancers:

```
/etc/init.d/wackamole start
```

Then run:

```
ifconfig
```

on both load balancers. The outputs should be different - one load balancer should now own the virtual IP address, e.g. like this:

```
lb1:~# ifconfig
eth0      Link encap:Ethernet  HWaddr 00:0C:29:A5:5B:93
          inet addr:192.168.0.100  Bcast:192.168.0.255  Mask:255.255.255.0
          inet6 addr: fe80::20c:29ff:fea5:5b93/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:9578 errors:0 dropped:0 overruns:0 frame:0
          TX packets:6347 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:8849468 (8.4 MiB)  TX bytes:811454 (792.4 KiB)
          Interrupt:177 Base address:0x1400

eth0:1    Link encap:Ethernet  HWaddr 00:0C:29:A5:5B:93
          inet addr:192.168.0.99  Bcast:192.168.0.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          Interrupt:177 Base address:0x1400

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          inet6 addr: ::1/128 Scope:Host
          UP LOOPBACK RUNNING  MTU:16436  Metric:1
          RX packets:51 errors:0 dropped:0 overruns:0 frame:0
          TX packets:51 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:3528 (3.4 KiB)  TX bytes:3528 (3.4 KiB)

lb1:~#
```

On the other load balancer, the output should be like this:

```
lb2:~# ifconfig
eth0      Link encap:Ethernet  HWaddr 00:0C:29:E0:78:92
          inet addr:192.168.0.101  Bcast:192.168.0.255  Mask:255.255.255.0
          inet6 addr: fe80::20c:29ff:fee0:7892/64  Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:6550 errors:0 dropped:0 overruns:0 frame:0
          TX packets:4109 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:8591676 (8.1 MiB)  TX bytes:377481 (368.6 KiB)
          Interrupt:177 Base address:0x1400

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          inet6 addr: ::1/128  Scope:Host
          UP LOOPBACK RUNNING  MTU:16436  Metric:1
          RX packets:11 errors:0 dropped:0 overruns:0 frame:0
          TX packets:11 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:824 (824.0 b)  TX bytes:824 (824.0 b)
```

```
lb2:~#
```

6 Starting HAProxy

Now we can start HAProxy:

[lb1/lb2:](#)

```
/etc/init.d/haproxy start
```

7 Testing

Our high-availability load balancer is now up and running.

You can now make HTTP requests to the virtual IP address `192.168.0.99` (or to any domain/hostname that is pointing to the virtual IP address), and you should get content from the backend web servers.

You can test its high-availability/failover capabilities by switching off one backend web server - the load balancer should then redirect all requests to the remaining backend web server. Afterwards, switch off the active load balancer (`lb1`) or stop Wackamole on the active load balancer - `lb2` should take over immediately. You can check that by running:

[lb2:](#)

```
ifconfig
```

You should now see the virtual IP address in the output on `lb2`:

```
lb2:~# ifconfig
eth0      Link encap:Ethernet  HWaddr 00:0c:29:e0:78:92
          inet addr:192.168.0.101  Bcast:192.168.0.255  Mask:255.255.255.0
          inet6 addr: fe80::20c:29ff:fee0:7892/64  Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:99050 errors:0 dropped:0 overruns:0 frame:0
          TX packets:56342 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:144164308 (137.4 MiB)  TX bytes:4502509 (4.2 MiB)
          Interrupt:177 Base address:0x1400

eth0:1    Link encap:Ethernet  HWaddr 00:0c:29:e0:78:92
          inet addr:192.168.0.99  Bcast:192.168.0.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          Interrupt:177 Base address:0x1400

lo        Link encap:Local Loopback
```

```
inet addr:127.0.0.1 Mask:255.0.0.0
inet6 addr: ::1/128 Scope:Host
UP LOOPBACK RUNNING MTU:16436 Metric:1
RX packets:60 errors:0 dropped:0 overruns:0 frame:0
TX packets:60 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:4296 (4.1 KiB) TX bytes:4296 (4.1 KiB)
```

```
lb2:~#
```

8 HAProxy Statistics

You might have noticed that we have used the options `stats enable` and `stats auth someuser:somepassword` in the HAProxy configuration in chapter 4. This allow us to access (password-protected) HAProxy statistics under the URL `http://192.168.0.99/haproxy?stats`. This is how it looks:

HAProxy version 1.3.15.2, released 2008/06/21
Statistics Report for pid 26686

> **General process information**

pid = 26686 (process #1, nbproc = 1)
 uptime = 0d 0h01m22s
 system limits : memmax = unlimited ; ulimit-n = 8205
 maxsock = 8205
 maxconn = 4096 (current conns = 1)

active UP
 active UP, going down
 active DOWN, going up
 active or backup DOWN

backup UP
 backup UP, going down
 backup DOWN, going up
 not checked

Note: UP with load-balancing disabled is reported as "NOLB".

Display option:
 • [Hide DOWN servers](#)
 • [Refresh now](#)
 • [CSV export](#)

webfarm

	Queue		Sessions				Bytes		Denied		Errors		Warnings			Server							
	Cur	Max	Limit	Cur	Max	Limit	Total	LbTot	In	Out	Req	Resp	Req	Conn	Resp	Retr	Redis	Status	Wght	Act	Bck	Chk	Dv
Frontend				1	1	2000	6		1861	2557	0	0	0					OPEN					
webA	0	0	-	0	1	-	2	2	717	1293		0		0	0	0	0	1m22s UP	1	Y	-	0	
webB	0	0	-	0	1	-	2	2	763	1002		0		0	0	0	0	1m22s UP	1	Y	-	0	
Backend	0	0		0	1	2000	4	4	1861	2557	0	0		0	0	0	0	1m22s UP	2	2	0		

If you don't need the statistics, just comment out or remove the *stats* lines from the HAProxy configuration.

9 Links

- HAProxy: <http://haproxy.1wt.eu>

- Wackamole: <http://www.backhand.org/wackamole/>
- Spread: <http://www.spread.org/>
- Debian: <http://www.debian.org>