

By Falko Timme

Published: 2007-05-29 17:42

Apache: Creating A Session-Aware Loadbalancer Using `mod_proxy_balancer` (Debian Etch)

Version 1.0

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

Last edited 05/26/2007

Since Apache 2.1, a new module called `mod_proxy_balancer` is available which lets you turn a system that has Apache installed into a loadbalancer. This loadbalancer retrieves requested pages from two or more backend webserver and delivers them to the user's computer. Users get the impression that they deal with just one server (the loadbalancer) when in fact there are multiple systems behind the loadbalancer that process the users' requests. By using a loadbalancer, you can lower the load average on your webserver. One important feature of `mod_proxy_balancer` is that it can keep track of sessions which means that a single user always deals with the same backend webserver. Most websites are database-driven nowadays with user logins etc., and you'd get weird results if a user logs in on one backend webserver, and then his next request goes to another backend webserver, meaning he'd get logged out again. You can avoid this by using `mod_proxy_balancer`'s session-awareness.

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


1 Preliminary Note

It should be noted that a loadbalancer can lower the load on your backend webserver, but it doesn't provide high-availability (as long as you use only one loadbalancer)! A single loadbalancer is a single point of failure (SPOF). If you need high-availability, you should use at least two loadbalancers (e.g. one as a hot-standby).

I will use three servers in this example: one loadbalancer with the address `www.example.com / example.com`, and two backend servers with the addresses `http1.example.com` and `http2.example.com`.

All three servers use Debian Etch as their operating system, and all three systems have Debian's Apache2 installed (its version is 2.2.3, so it comes with `mod_proxy_balancer` by default). On `http1.example.com` and `http2.example.com` I have installed a database-driven web application: [phpBB2](#), a famous forum software. Both phpBB2 installations are identical and use the same database. For debugging purposes, I've built-in a small difference in both

installations: if the page is delivered by `http1.example.com`, it shows `http1.example.com` in the header:



And if it's delivered by `http2.example.com`, it shows `http2.example.com`:



That way I can control if sessions are handled correctly by `mod_proxy_balancer`. Of course, on a production system both pages would be the same.

Session-tracking is a bit tricky in `mod_proxy_balancer` because it expects a certain cookie format, and the name of the session variable can differ from application to application. Fortunately I've found this page:

http://www.markround.com/archives/33-Apache-mod_proxy-balancing-with-PHP-sticky-sessions.html which has a great and easy solution for this problem that I'm going to use here.

2 Preparing The Backend Servers

First we must prepare our backend webserver `http1.example.com` and `http2.example.com`. We enable `mod_rewrite` like this:

[http1.example.com / http2.example.com:](http1.example.com/)

```
a2enmod rewrite
```

```
/etc/init.d/apache2 force-reload
```

Then we open the Apache vhost configuration of our phpBB2 site on `http1.example.com` and add the following lines to it:

<http1.example.com>:

```
[...]  
RewriteEngine On  
RewriteRule .* - [CO=BALANCEID:balancer.http1:.example.com]  
[...]
```

(Make sure that you replace `http1` and `.example.com` according to your needs!)

Restart Apache afterwards:

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

Now we do the same on `http2.example.com`:

<http2.example.com>:

```
[...]  
RewriteEngine On  
RewriteRule .* - [CO=BALANCEID:balancer.http2:.example.com]  
[...]
```

(Make sure that you replace *http2* and *.example.com* according to your needs!)

Restart Apache afterwards:

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

That's all we have to configure on the backend servers.

3 Configuring The Loadbalancer

On our loadbalancer *www.example.com* we must enable a few Apache modules:

www.example.com:

```
a2enmod proxy  
  
a2enmod proxy_balancer  
  
a2enmod proxy_http  
  
a2enmod status
```

and then restart Apache:

```
/etc/init.d/apache2 force-reload
```

I'm assuming that we don't run any other websites on the loadbalancer, so we can use Debian's default document root */var/www* for our loadbalancer.

mod_proxy_balancer comes with a "Balancer Manager", a small web interface where you can tweak a few settings. We create the directory */var/www/balancer-manager* for it which we password-protect so that only we have access to it:

```
mkdir /var/www/balancer-manager

htpasswd -c /var/.htpasswd admin
```

(You can replace *admin* with any username you like.)

```
vi /var/www/balancer-manager/.htaccess
```

```
AuthType Basic
AuthName "Members Only"
AuthUserFile /var/.htpasswd
<limit GET PUT POST>
require valid-user
</limit>
```

Now we come to the vhost configuration of the loadbalancer. Apache's default vhost configuration on Debian Etch is located in */etc/apache2/sites-available/default*, so we replace that with our own configuration:

```
cp /etc/apache2/sites-available/default /etc/apache2/sites-available/default_orig

cat /dev/null > /etc/apache2/sites-available/default

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

```
NameVirtualHost *
<VirtualHost *>
    ServerName www.example.com
    ServerAlias example.com
    DocumentRoot /var/www/
    ProxyRequests Off

    <Proxy *>
        Order deny,allow
        Allow from all
    </Proxy>

    ProxyPass /balancer-manager !
    ProxyPass / balancer://mycluster/ stickysession=BALANCEID nofailover=On
    ProxyPassReverse / http://http1.example.com/
    ProxyPassReverse / http://http2.example.com/
    <Proxy balancer://mycluster>
        BalancerMember http://http1.example.com route=http1
        BalancerMember http://http2.example.com route=http2
        ProxySet lbmethod=byrequests
    </Proxy>

    <Location /balancer-manager>
        SetHandler balancer-manager

        Order deny,allow
        Allow from all
    </Location>
</VirtualHost>
```

It is very important that you set all slashes (/) EXACTLY as shown in this example, especially the trailing slashes!

Please make sure that the value of *stickysession* is the name of the cookie (*BALANCEID*) in our rewrite rules from chapter 2. Also, the values of *route* in the *BalancerMember* lines must be the respective value that we set in the rewrite rules (*http1* or *http2*). The *stickysession* and *route* values take care of the correct session handling of *mod_proxy_balancer*.

It's also important that you have as many *ProxyPassReverse* lines as you have *BalancerMember* lines (one *ProxyPassReverse* line for each *BalancerMember*).

The *ProxyPass /balancer-manager !* line makes sure that requests to *www.example.com/balancer-manager* aren't routed to the backend servers.

You can find more details about all settings and further fine-tuning options on http://httpd.apache.org/docs/2.2/mod/mod_proxy.html#proxypass and http://httpd.apache.org/docs/2.2/mod/mod_proxy_balancer.html.

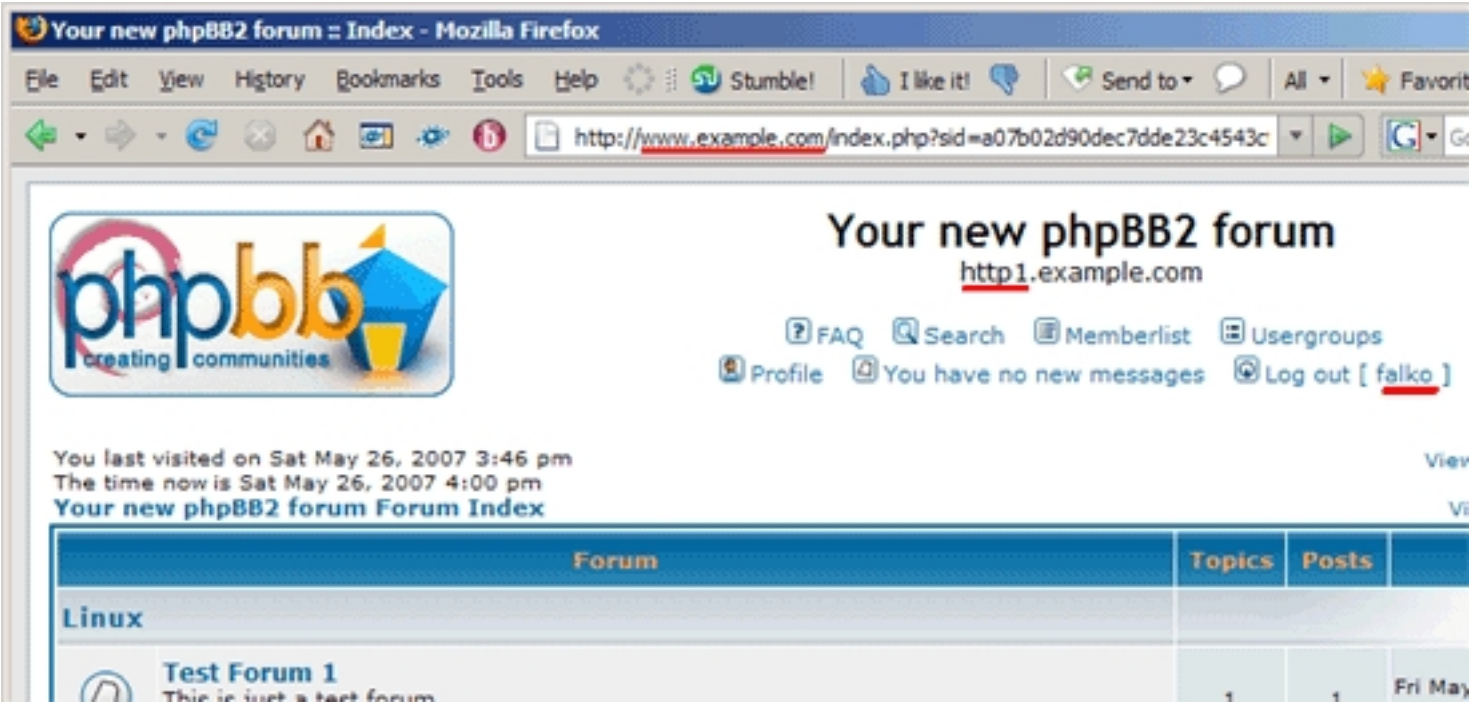
Now that we've finished the configuration, we must restart Apache:

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

4 Testing Our Setup

Our setup is now ready to be used. To test if sessions are handled correctly, I open two browsers (because of the cookies), e.g. Firefox and Internet Explorer, and go to *http://www.example.com* or *http://example.com*. In both cases, you should see the phpBB2 forum, delivered by one of the backend servers (which you don't see, you should always have *www.example.com* or just *example.com* in the browser's address bar).

It's possible that the pages in both browsers are delivered by the same backend server - if that's the case, delete the cookies in your browsers or open another browser (Opera, Seamonkey, ...) and try again until you see that the pages in your two different browsers come from different backend servers (I can see it because I changed the header, as described in chapter 1). Then log in with two different users (e.g. *falko* and *till*) that you have created before and browse the forum with these two different users. In my case *falko*'s content is always delivered by *http1.example.com*:



Whereas *till*'s content comes from *http2.example.com*:



So sessions are handled correctly by `mod_proxy_balancer`!

5 The Balancer Manager

Direct your browser to `http://www.example.com/balancer-manager` or `http://example.com/balancer-manager`. You will be prompted for a username and password (the username and password you've created in chapter 3).

After the login you will see a simple web interface for managing the loadbalancer:



Balancer Manager - Mozilla Firefox

File Edit View History Bookmarks Tools Help Stumble! I like it! Send to All Favorites Friends Tools

http://www.example.com/balancer-manager

Load Balancer Manager for www.example.com

Server Version: Apache/2.2.3 (Debian) PHP/5.2.0-8+etch4 mod_ssl/2.2.3 OpenSSL/0.9.8c
 Server Built: Mar 27 2007 15:06:55

LoadBalancer Status for [balancer://mycluster](#)

StickySession Timeout FailoverAttempts Method
 BALANCEID 0 1 byrequests

Worker URL	Route	RouteRedir	Factor	Status
http://http1.example.com	http1		1	Ok
http://http2.example.com	http2		1	Ok

Apache/2.2.3 (Debian) PHP/5.2.0-8+etch4 mod_ssl/2.2.3 OpenSSL/0.9.8c Server at www.example.com Port 80

Copyright © 2007 All Rights Reserved.

Load Balancer Manager for www.example.com

Server Version: Apache/2.2.3 (Debian) PHP/5.2.0-8+etch4 mod_ssl/2.2.3 OpenSSL/0.9.8c
Server Built: Mar 27 2007 15:06:55

LoadBalancer Status for [balancer://mycluster](#)

StickySession	Timeout	FailoverAttempts	Method	BALANCEID	
	0	1	byrequests		

Worker URL	Route	RouteRedir	Factor	Status
http://http1.example.com	http1		1	Ok
http://http2.example.com	http2		1	Ok

Edit worker settings for <http://http1.example.com>

Load factor:

Route:

Route Redirect:

Status: Disabled: ☐ | Enabled: ☒

Apache/2.2.3 (Debian) PHP/5.2.0-8+etch4 mod_ssl/2.2.3 OpenSSL/0.9.8c Server at www.example.com Port 80

You can find an overview of possible configuration options on http://httpd.apache.org/docs/2.2/mod/mod_proxy.html#proxypass.

6 Links

- Apache: <http://httpd.apache.org>
- mod_proxy: http://httpd.apache.org/docs/2.2/mod/mod_proxy.html
- mod_proxy_balancer: http://httpd.apache.org/docs/2.2/mod/mod_proxy_balancer.html
- Apache mod_proxy balancing with PHP sticky sessions:
http://www.markround.com/archives/33-Apache-mod_proxy-balancing-with-PHP-sticky-sessions.html
- Debian: <http://www.debian.org>