

Load Balancing Xpand with HAProxy

Xpand requires an external tcp load balancer to evenly distribute load across the cluster. Xpand recommends using [HAProxy](#) or the [Xpand Monitor](#).

Instructions to install HAProxy:

1) Install HAProxy to a server that will be used for load balancing and routing traffic to the cluster. If you are running Centos or Redhat you can install HAProxy with yum by running the following command from bash:

```
shell> sudo yum install haproxy
```

If your server doesn't support YUM, you can [download HAProxy here](#).

2) Once you have HAProxy installed, you will need to edit the haproxy.cfg file located at: /etc/haproxy/haproxy.cfg. A sample haproxy.cfg file is included below. You will need to change the timeout variables in the defaults section and then set two listen sections, the first for Xpand access, and the second for XpandGUI.

The first thing to do in the haproxy.cfg file is to change the timeout variables in the default section. You want to increase the timeout connect, timeout client, and timeout server to 1d. This is because you'll want Xpand to handle the timeout settings instead of HAProxy. There are several other timeout variables provided by HAProxy that can be set as you see fit. The next thing is to change httplog to tcplog, as we are going to be sending mostly tcp traffic to the cluster.

Example:

```
defaults
  log                  global
  option               tcplog
  option               dontlognull
  option               redispatch
  retries              3
  timeout connect     1d
  timeout client      1d
  timeout server      1d
  maxconn             3000
```



If you have queries that run for longer than 1 day, you should consider upping the timeout values even higher.

In the first listen section, you will need to supply the name, in this case, Xpand, and the IP address and port of the VIP to which clients will connect. This VIP can be an additional IP configured on the box running HAProxy. You will want to use the port from which you plan to access Xpand.

Now you will need to set the mode to mode tcp as you will be balancing tcp and not http traffic.

Next, you will need to select what type of balance you want HAProxy to perform. The suggested methods are leastconn and roundrobin. You can find out more about the balance modes by searching for "balance <algorithm>" in the Configuration Manual of the [HAProxy documentation](#) that corresponds to your installed version.

For each node to which you will be balancing traffic, you will need a server entry. For each server, you will need to supply the name and IP or DNS hostname of the node. Finally, for each server entry, you will want to add the check port argument, and have it set to 3581, the Xpand health check port. Setting this check port 3581 argument will ensure that traffic is only balanced to nodes that are in quorum.

Example:

```
listen ClustrixDB 10.1.1.100:3306
  mode tcp
  balance leastconn
  server examplenode001 10.1.1.101:3306 check port 3581
  server examplenode002 10.1.1.102:3306 check port 3581
  server examplenode003 10.1.1.103:3306 check port 3581
```

The second listen section handles configuration for load balancing XpandGUI in the case of a node failure. First, you supply the name gui and the IP address from which you are going to be accessing the XpandGUI. Also, make sure to include port 80.

Now make sure the mode is set to http with mode http.

Next, you will need to select what type of balance you want HAProxy to perform. The suggested methods are leastconn and roundrobin.

For each node to which you will be balancing web traffic, you will need a server entry. For each server, you will need to supply the name followed by _gui and IP or DNS hostname of the node.

Example:

```
listen gui 10.1.1.100:80
  mode http
  balance leastconn
  server examplenode001_gui 10.1.1.101:80
  server examplenode002_gui 10.1.1.102:80
  server examplenode003_gui 10.1.1.103:80
```

3) Once you have all the settings correct, you will need to restart HAProxy to make them take effect.

```
shell> sudo service haproxy
restart
```

Below is the full haproxy.cfg file used for internal testing by the Xpand Support team.

```
haproxy.cfg
#-----
# Global settings
#-----
global
  log          127.0.0.1 local2
  chroot       /var/lib/haproxy
  pidfile      /var/run/haproxy.pid
  maxconn      4000
  user         haproxy
  group        haproxy
  daemon
  # turn on stats unix socket
  stats socket /var/lib/haproxy/stats

#-----
# common defaults that all the 'listen' and 'backend' sections will
# use if not designated in their block
#-----
defaults
  log            global
  option         tcplog
  option         dontlognull
  option         redispatch
  retries        3
  timeout connect 1d
  timeout client 1d
  timeout server 1d
  maxconn        3000

listen ClustrixDB 10.1.1.100:3306
  mode tcp
  balance leastconn
  server examplenode001 10.1.1.101:3306 check port 3581
  server examplenode002 10.1.1.102:3306 check port 3581
  server examplenode003 10.1.1.103:3306 check port 3581

listen gui 10.1.1.100:80
  mode http
  balance leastconn
  server examplenode001_gui 10.1.1.101:80
  server examplenode002_gui 10.1.1.102:80
  server examplenode003_gui 10.1.1.103:80
```

For more information on configuring HAProxy, please see their [documentation](#).