

# carp(4) ip loadbalancing (patch included)

---

*Source:* <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/net/2008-01/msg00236.html>

---

- *From:* "Felix J. Ogris" <[fjo-lists@xxxxxxxx](mailto:fjo-lists@xxxxxxxx)>
  - *Date:* Thu, 24 Jan 2008 19:08:25 +0100
- 

Hi,

I have extended ip\_carp.c to provide loadbalancing on an ip basis, eg. to setup an active/active cluster. The algorithm is quite simple. Each cluster consists of N nodes. If an IPv4/IPv6 packet reaches node X, then it evaluates  $X == N \bmod \text{source\_address\_of\_IP\_packet}$ . If this is true, then the node will process that packet (carp\_forus() returns 1). I had to introduce a new carp type (#define CARP\_HELO 0x00) which by default is multicasted every 3 seconds (advertising base). Each node maintains a list of its IPv4 and IPv6 neighbours. Neighbours are timed out after  $3 * \text{advbase}$  seconds. Network layer protocols other than IPv4/IPv6 are handled by the carp master. TCP packets, which have their ACK flag set or their SYN flag not set, are checked for an entry in the syncache or in tcbinfo. If an entry is found, then that TCP packet will not be handled by the current node (carp\_forus() returns 0). This guarantees that an established TCP connection will stick to a host, even if a dead node reenters the cluster.

I have put my patch (against 6.3-RELEASE) under the same license as the FreeBSD kernel itself. So feel free to apply it to the official tree – or to punish me for my lousy work :-)

Regards,  
Felix

***Attachment:*** [carp\\_aa.patch](#)

*Description:* Binary data

---

freebsd-net@xxxxxxxxxxxxx mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-net>

To unsubscribe, send any mail to "freebsd-net-unsubscribe@xxxxxxxxxxxxx"