

Table of Contents

Show Status 1

Show Status

```
cat /proc/net/bonding/bond0
```

Number	Mode	Description
0	balance-rr	Transmits packets through network interface cards in sequential order. Packets are transmitted in a loop that begins with the first available network interface card in the bond and end with the last available network interface card in the bond. All subsequent loops then start with the first available network interface card. Mode 0 offers fault tolerance and balances the load across all network interface cards in the bond. NOTE - Mode 0 cannot be used in conjunction with bridges, and is therefore not compatible with virtual machine logical networks. This mode provides load balancing and fault tolerance.
1	active-backup	Only one slave in the bond is active. A different slave becomes active if, and only if, the active slave fails. The bond's MAC address is externally visible on only one port (network adapter) to avoid confusing the switch. In bonding version 2.6.2 or later, when a failover occurs in active-backup mode, bonding will issue one or more gratuitous ARPs on the newly active slave. One gratuitous ARP is issued for the bonding master interface and each VLAN interfaces configured above it, provided that the interface has at least one IP address configured. Gratuitous ARPs issued for VLAN interfaces are tagged with the appropriate VLAN id. This mode provides fault tolerance.
2	balance-xor	Transmit based on the selected transmit hash policy. The default policy is a simple [(source MAC address XOR'd with destination MAC address) modulo slave count]. Alternate transmit policies may be selected via the xmit_hash_policy option. This mode provides load balancing and fault tolerance.
3	broadcast	Transmits all packets to all network interface cards. This mode provides fault tolerance.
4	802.3ad	Creates aggregation groups that share the same speed and duplex settings. Utilizes all slaves in the active aggregator according to the 802.3ad specification. Slave selection for outgoing traffic is done according to the transmit hash policy, which may be changed from the default simple XOR policy via the xmit_hash_policy option, documented below. Note that not all transmit policies may be 802.3ad compliant, particularly in regards to the packet mis-ordering requirements of section 43.2.4 of the 802.3ad standard. Differing peer implementations will have varying tolerances for noncompliance. A switch that supports IEEE 802.3ad Dynamic link aggregation. Most switches will require some type of configuration to enable 802.3ad mode.
5	balance-tlb	Adaptive transmit load balancing: Channel bonding that does not require any special switch support. The outgoing traffic is distributed according to the current load (computed relative to the speed) on each slave. Incoming traffic is received by the current slave. If the receiving slave fails, another slave takes over the MAC address of the failed receiving slave. NOTE - Mode 5 cannot be used in conjunction with bridges, therefore it is not compatible with virtual machine logical networks.

Number	Mode	Description
6	balance-alb	<p>Adaptive load balancing: Includes balance-tlb plus receive load balancing (rlb) for IPV4 traffic, and does not require any special switch support. The receive load balancing is achieved by ARP negotiation. The bonding driver intercepts the ARP Replies sent by the local system on their way out and overwrites the source hardware address with the unique hardware address of one of the slaves in the bond such that different peers use different hardware addresses for the server. Receive traffic from connections created by the server is also balanced. When the local system sends an ARP Request the bonding driver copies and saves the peer's IP information from the ARP packet. When the ARP Reply arrives from the peer, its hardware address is retrieved and the bonding driver initiates an ARP reply to this peer assigning it to one of the slaves in the bond. A problematic outcome of using ARP negotiation for balancing is that each time that an ARP request is broadcast it uses the hardware address of the bond. Hence, peers learn the hardware address of the bond and the balancing of receive traffic collapses to the current slave. This is handled by sending updates (ARP Replies) to all the peers with their individually assigned hardware address such that the traffic is redistributed. Receive traffic is also redistributed when a new slave is added to the bond and when an inactive slave is re-activated. The receive load is distributed sequentially (round robin) among the group of highest speed slaves in the bond. When a link is reconnected or a new slave joins the bond the receive traffic is redistributed among all active slaves in the bond by initiating ARP Replies with the selected mac address to each of the clients. The updelay parameter (detailed below) must be set to a value equal or greater than the switch's forwarding delay so that the ARP Replies sent to the peers will not be blocked by the switch. Base driver support for setting the hardware address of a device while it is open. This is required so that there will always be one slave in the team using the bond hardware address (the curr_active_slave) while having a unique hardware address for each slave in the bond. If the curr_active_slave fails its hardware address is swapped with the new curr_active_slave that was chosen. NOTE - Mode 6 cannot be used in conjunction with bridges, therefore it is not compatible with virtual machine logical networks.</p>

From:
<https://wiki.janforman.com/> - **wiki.janforman.com**

Permanent link:
<https://wiki.janforman.com/linux:bond>

Last update: **2018/02/02 22:24**

