

# ALOHA LOAD BALANCER ACTIVE-ACTIVE BONDING

## "APPNOTE" #0006 — CONFIGURING ACTIVE-ACTIVE BONDING

*This application note is intended to help you configure bonding in order to ensure high availability of your links in active-active mode within the ALOHA Load Balancer solution.*

### CONSTRAINTS

"Port trunking" mode must be supported and configured on the network switch ports the Aloha devices are connected to.

### OBJECTIVE

Dynamically aggregate the Aloha network interfaces in order to increase bandwidth and manage fault tolerance should one of them fail.

### COMPLEXITY



### CHANGELOG


2013-12-05: xmit\_hash\_policy configuration changes

2012-12-13: bonding modes description, advanced configuration to set-up /sys parameters

2011-10-19: update for Aloha 4.x and above

2009-03-28: initial version

## TARGET NETWORK DIAGRAM



## ALOHA VERSION 4.0 AND ABOVE

### BONDING CONFIGURATION

When setting up a bonding interface, a few options are available:

- `slave <iface> <iface>...` : interfaces to add in the bonding
- `bonding primary <iface>` : primary slave for current bond
- `bonding mode <mode>` : bonding mode
- `bonding updelay <ms>` : bonding delay check in ms (default: 1000)
- `bonding downdelay <ms>` : bonding delay check in ms (default: 1000)

Bonding mode can have different values:

- **0: balance-rr mode:** output interfaces are chosen in a round-robin method
- **1: active-backup mode:** failover mode, one interface is active at a time
- **2: balance-xor mode:** Transmit based on [(source MAC address XOR'd with destination MAC address) modulo slave count]. This selects the same slave for each destination MAC address. This mode provides load balancing and fault tolerance and does not requires any configuration on the switch.
- **4: 802.3ad:** IEEE 802.3ad (**LACP**) Dynamic link aggregation. The switch must be compatible and properly configured
- **5: balance-tlb:** "transmit load-balancing", outgoing interface choice is made based on interfaces load
- **6: balance-alb:** "adaptive load-balancing", balance-tlb mode with a reception load-balancing mode made using ARP protocol

## EXAMPLE

You can directly configure the **network** service from the **Services** tab of the GUI.

### **Before modifications:**

```
service network eth0
##### Interface Interface eth0: .
auto on
mtu 1500
ip address 10.0.32.10/16
ip route default 10.0.32.1

service network eth1
##### Interface eth1: .
auto on
mtu 1500
```

### **After modifications:**

```
service network eth0
##### Interface Interface eth0: .
auto on
mtu 1500

service network eth1
##### Interface eth1: .
auto on
mtu 1500

service network bond0
##### Interface bond0: active/backup aggregate of eth0 and eth1
slave eth0 eth1
bonding primary eth0
bonding mode 2
bonding updelay 500
bonding downdelay 500
ip address 10.0.32.10/16
ip route default 10.0.32.1
mtu 1500
```



**Warning: the network interfaces used in the bonding configuration must no longer have a specific individual configuration!**

## FOR ADVANCED USERS ONLY

It's possible to change **bonding** driver parameters available through **/sys**.

You just need to append the parameter name and its value to the configuration of the bonding interface.

In exemple, to choose the output interface in a LACP aggregation from a hash of the destination IP and Port (more efficient than the single hash on destination MAC address), you must setup the xmit-hash-policy parameter to 1:

```
service network bond0
##### Interface bond0: XOR policy of eth0 and eth1
slave eth0 eth1
bonding primary eth0
bonding mode 2
bonding updelay 500
bonding downdelay 500
ip address 10.0.32.10/16
ip route 10.0.32.1
mtu 1500
bonding xmit_hash_policy 1
```

## ALOHA VERSION 3.X ONLY

### EXTRACT OF THE CONFIGURATION OF THE "NETWORK" SERVICE

#### **Before modifications:**

```
service network eth0
##### Interface Interface eth0: .
auto on
mtu 1500
ip address 10.0.32.10/16
ip route default 10.0.32.1

service network eth1
##### Interface eth1: .
auto on
mtu 1500
```

#### **After modifications:**

```
service network eth0
##### Interface Interface eth0: .
auto on
mtu 1500

service network eth1
##### Interface eth1: .
auto on
mtu 1500

service network bond0
##### Interface bond0: XOR policy of eth0 and eth1
slave eth0 eth1
modprobe bond0 mode=balance-xor updelay=1000 downdelay=1000
ip address 10.0.32.10/16
ip route default 10.0.32.1
mtu 1500
```

You can directly configure the **network** service from the **Services** tab.

To meet this objective, simply specify the network interfaces where bonding must be enabled (slave eth0 eth1).

Then provide the name of the new interface (bond0 in this example) and specify the bonding operating mode; this example uses "balance-xor" mode.

Delays for checking the status of the interfaces are specified using the "updelay" and "downdelay" parameters (values are specified in milliseconds).

Finally, configure the IP address and the route for the new interface.



**Warning: the network interfaces used in the bonding configuration must no longer have an individual configuration!**

## VALIDATION OF THE CONFIGURATION

network		Auto.					need restart
eth0	Auto.						
eth1	Auto.						
eth2	Auto.						
bond0	Auto.						

If successful, a new interface called **bond0** is created.

Just click on the **network restart icon**  to tell the Aloha to apply the new configuration.

## BASIC TROUBLESHOOTING

### Diagnostic Tools

interfaces (ignored, ignored, ignored) Retries: 5 Timeout: 5s  
\* Optional

Diagnostic Tool ( )

```
# /bin/netstat -nie
Kernel Interface table
bond0    Link encap:Ethernet  HWaddr 00:0D:B9:13:52:2D
         inet addr:10.0.32.10  Bcast:0.0.0.0  Mask:255.255.0.0
         UP BROADCAST RUNNING MASTER MULTICAST  MTU:1500  Metric:1
         RX packets:0 errors:0 dropped:0 overruns:0 frame:0
         TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
         collisions:0 txqueuelen:0
         RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)

eth0     Link encap:Ethernet  HWaddr 00:0D:B9:13:52:2D
         UP BROADCAST NOARP SLAVE MULTICAST  MTU:1500  Metric:1
         RX packets:0 errors:0 dropped:0 overruns:0 frame:0
         TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
         collisions:0 txqueuelen:1000
         RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)
         Interrupt:11 Base address:0x2000

eth1     Link encap:Ethernet  HWaddr 00:0D:B9:13:52:2D
         UP BROADCAST NOARP SLAVE MULTICAST  MTU:1500  Metric:1
         RX packets:0 errors:0 dropped:0 overruns:0 frame:0
         TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
         collisions:0 txqueuelen:1000
         RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)
         Interrupt:12 Base address:0x4000
```

In order to ensure that the bonding is implemented properly, use the built-in **interfaces** tool available in the **DiagTools** tab. If the configuration is correct, the "bond0" interface should appear with the "MASTER MULTICAST" label and both of the network interfaces "eth0" and "eth1" should be labeled as "SLAVE MULTICAST" and have the same MAC address as the "bond0" interface. Furthermore, each interface should be "UP".