

SimplStor

NETWORK STORAGE SOLUTIONS BY



White Paper

Network Bonding on SimplStor running RHEL7/CentOS7

Date: May 24, 2017



Why Bond Network Ports?

Channel bonding is a method of aggregating multiple network interfaces into a single logical interface, enabling two or more physical ports to act as one. You enjoy the benefits of load balancing the network interfaces and implement a more fault tolerant environment by virtue of the fact that if one ethernet port fails the other(s) will continue to function with no outage to your user community.

Important Note

All the commands described in this paper must be run as the “root” user. Because you will be modifying the network interfaces, it is highly advised that these commands be run from a console session. Before proceeding with these instructions, make backup copies of all the “ifcfg” files in the /etc/sysconfig/network-scripts directory so that you can revert to a non-bonded configuration if necessary.

This example only shows two bonded ports. If you desire to bond three or more, just follow the example and define them as your requirements dictate. For a thorough description of networking administration on RHEL7/CentOS7 please consult the “Red Hat Enterprise Linux 7 Networking Guide” available online at https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Networking_Guide/index.html

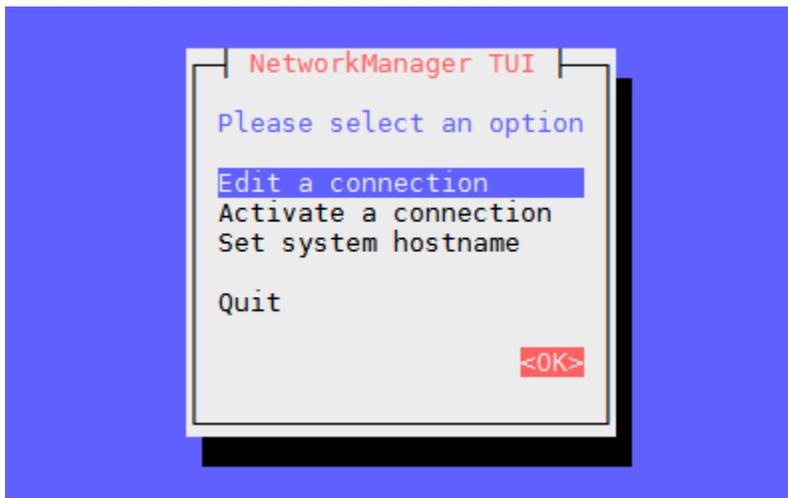
Channel Bonding

Manually editing the interface configuration files can be error prone, so the method described here makes use of “nmtui” – the Network Manager Text User Interface. This interface presents all the information you’ll need to configure the appropriate “ifcfg” scripts that are used by the system. Prior to beginning, ensure that you have physically connected the network cabling to appropriate network switch ports in your environment. Run the “ip link show” command to gather the interface names you will need for bonding.

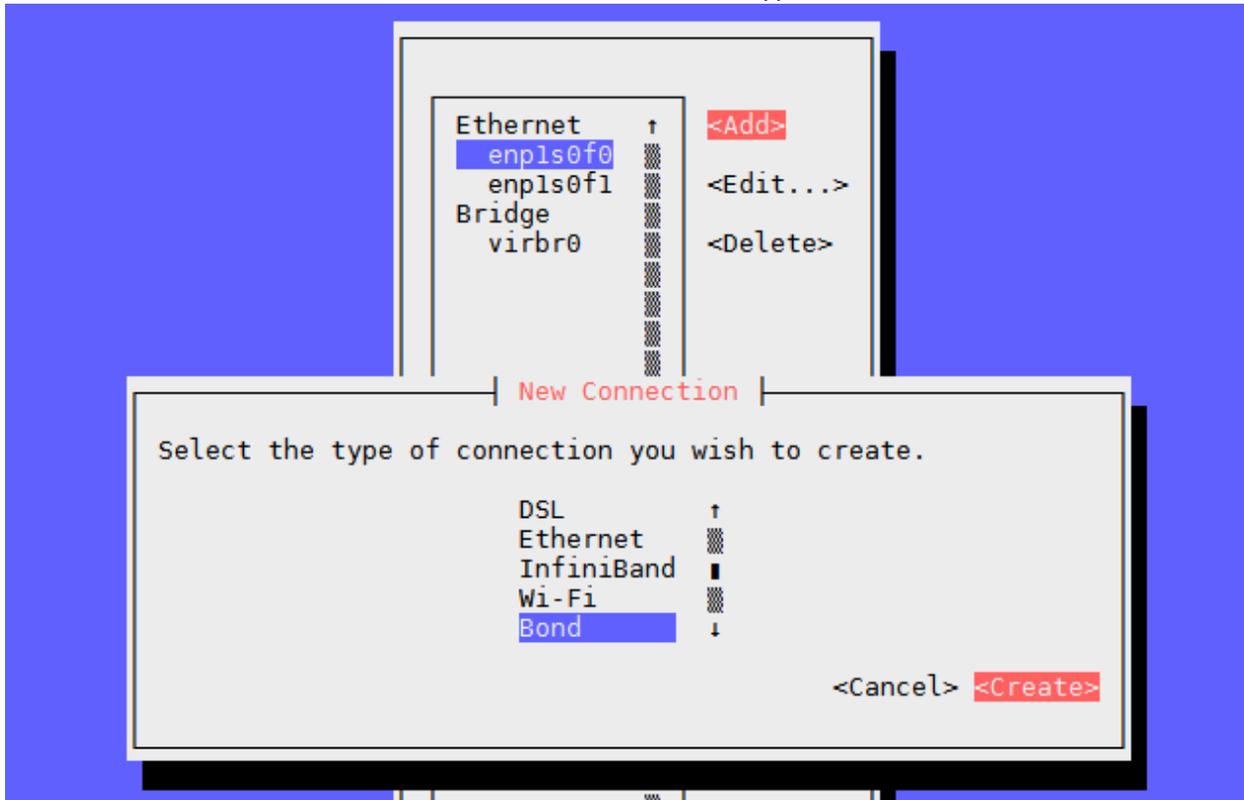
```
[root@simplstor ~]# ip link show
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode
DEFAULT qlen 1
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: enp1s0f0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP
mode DEFAULT qlen 1000
    link/ether 00:25:90:30:xx:yy brd ff:ff:ff:ff:ff:ff
3: enp1s0f1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP
mode DEFAULT qlen 1000
    link/ether 00:25:90:30:xx:yz brd ff:ff:ff:ff:ff:ff
4: virbr0: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc noqueue state
DOWN mode DEFAULT qlen 1000
    link/ether 52:54:00:4a:ww:vv brd ff:ff:ff:ff:ff:ff
5: virbr0-nic: <BROADCAST,MULTICAST> mtu 1500 qdisc pfifo_fast master virbr0
state DOWN mode DEFAULT qlen 1000
    link/ether 52:54:00:4a:ww:vv brd ff:ff:ff:ff:ff:ff
```

The type or mode of network bonding you choose may be dependent on the features that are supported by your switch. Please consult with the RHEL 7 Networking Guide and your switch documentation for a complete description of modes that are available to you.

Once you have the current interface names and have decided on the bonding mode, run “nmtui” and select “Edit a connection”.



Select <Add>, then select “Bond” for the new connection type and <Create>.



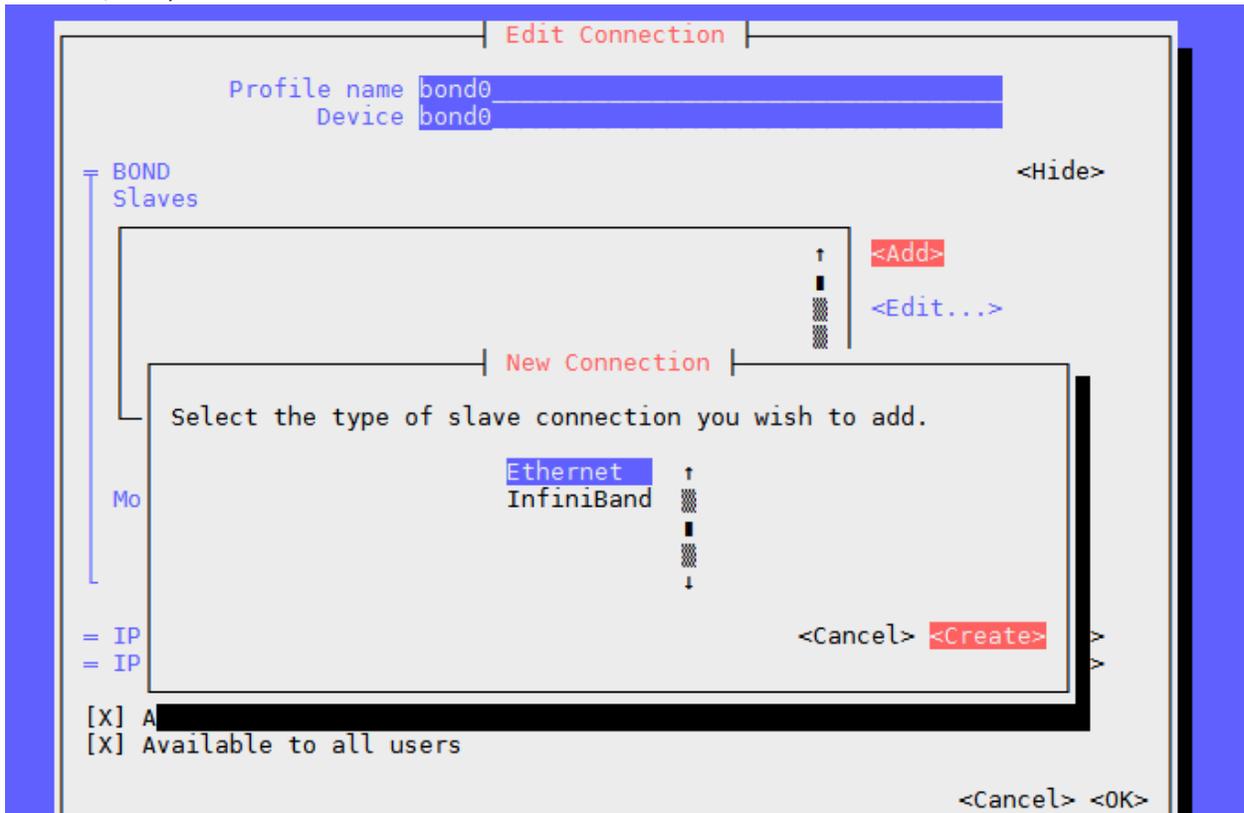
SimplStor Division:

sales@simplstor.com

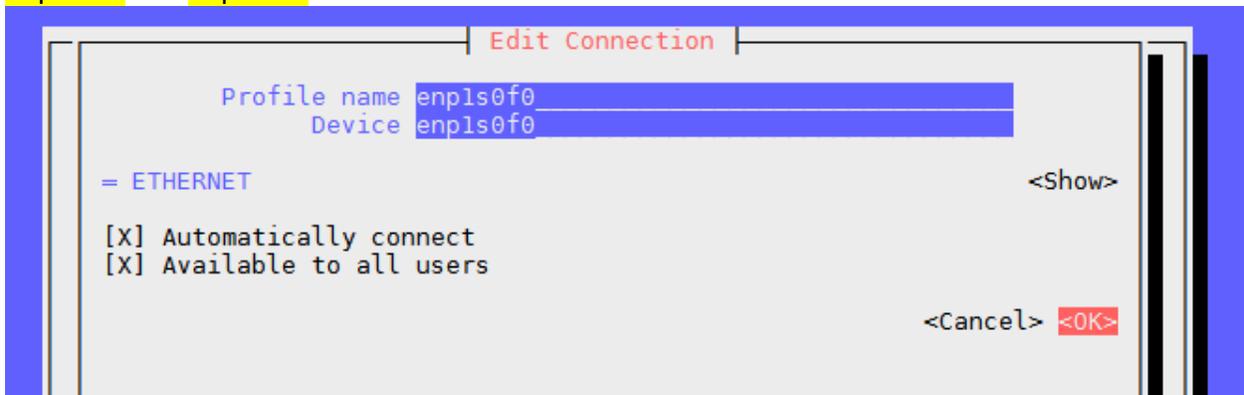
support@simplstor.com

All rights reserved. The information contained herein is subject to change without notice. Reproduction, adaptation, or translation without prior written permission is prohibited, except as allowed under the copyright laws. SimplStor is a trademark of Zerowait Corporation. The Zerowait Logo, ZHA, and Zerowait are registered trademarks of the Zerowait Corporation. ©2017 Zerowait Corporation. Page 4

Enter the Profile name and Device name (traditionally these have been named 'bond0', 'bond1', etc.) and then <Add> "slave" Ethernet interfaces.



Using the interface names you collected earlier, add these physical interfaces as the slaves for your new bond. For most SimplStor systems the two on-board Ethernet interfaces are named **enp1s0f0** and **enp1s0f1**.



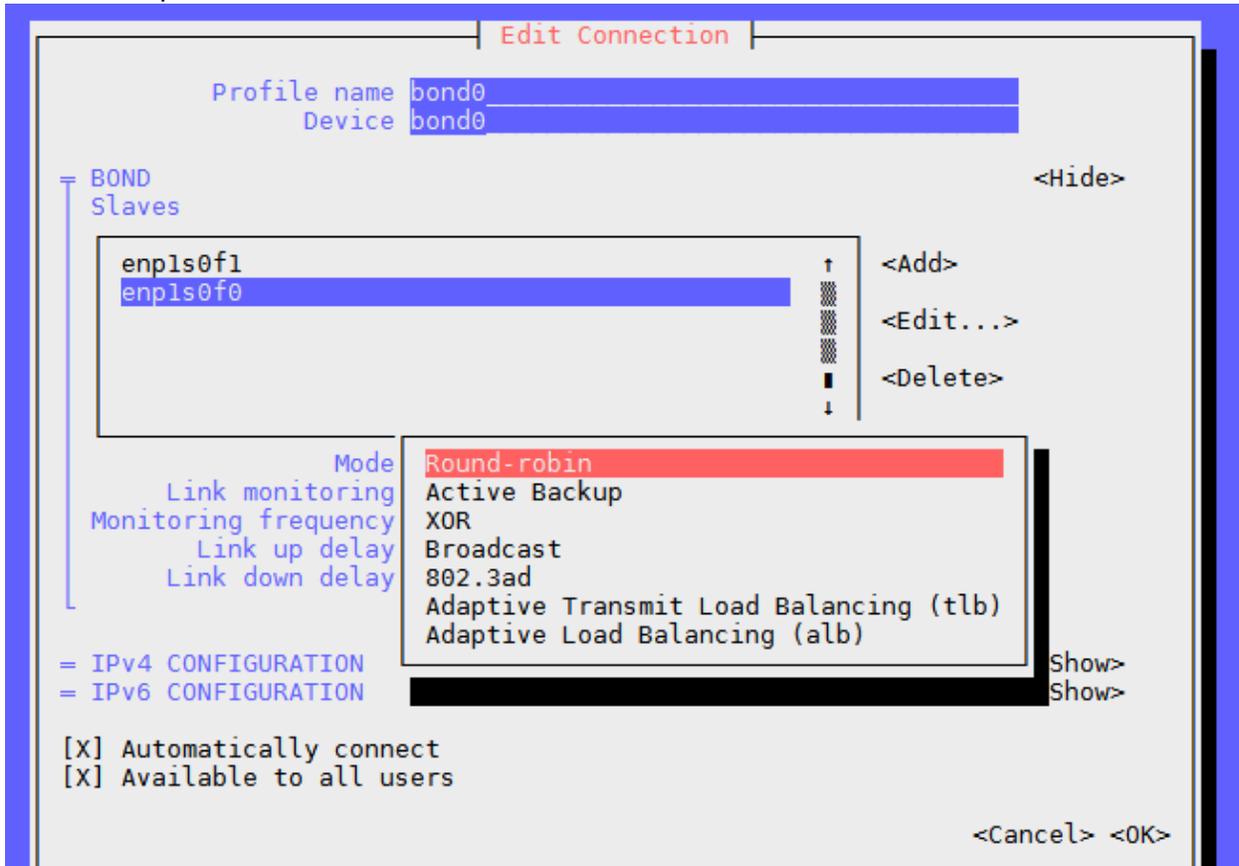
SimplStor Division:

sales@simplstor.com

support@simplstor.com

After adding the slaves, select the bonding Mode. The first three listed (Round-robin, Active Backup and XOR) can be used with any switch. The other Modes may require device driver support and proper configuration of the switch. There is plenty of discussion of these modes available in the Internet to help you select the best one for your requirements.

In our example we choose Round-robin.



Available Modes:

- Round-robin (0) — Sets a round-robin policy for fault tolerance and load balancing. Transmissions are received and sent out sequentially on each bonded slave interface beginning with the first one available.
- Active backup (1) — Sets an active-backup policy for fault tolerance. Transmissions are received and sent out via the first available bonded slave interface. Another bonded slave interface is only used if the active bonded slave interface fails.
- XOR (2) — Sets an XOR (exclusive-or) policy for fault tolerance and load balancing. Using this method, the interface matches up the incoming request's MAC address with the

SimplStor Division:

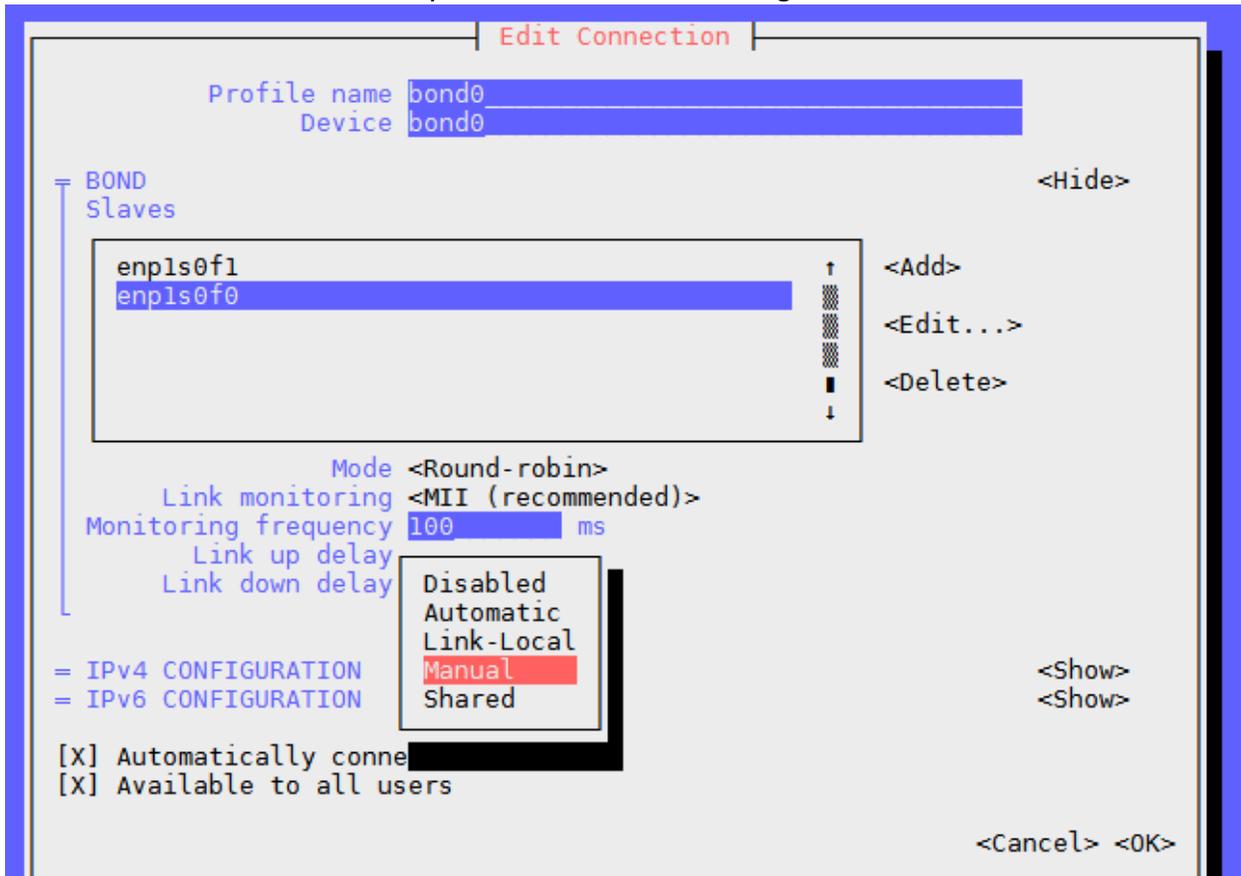
sales@simplstor.com

support@simplstor.com

MAC address for one of the slave NICs. Once this link is established, transmissions are sent out sequentially beginning with the first available interface.

- Broadcast (3) — Sets a broadcast policy for fault tolerance. All transmissions are sent on all slave interfaces.
- 802.3ad (4) — Sets an IEEE 802.3ad dynamic link aggregation policy. Creates aggregation groups that share the same speed and duplex settings. Transmits and receives on all slaves in the active aggregator. Requires a switch that is 802.3ad compliant.
- Adaptive Transmit Load Balance-tlb (5) — Sets a Transmit Load Balancing (TLB) policy for fault tolerance and load balancing. The outgoing traffic is distributed according to the current load on each slave interface. Requires Ethtool support in the base drivers and should be supported by all SimplStor network cards.
- Adaptive Load Balance-alb (6) — Sets an Active Load Balancing (ALB) policy for fault tolerance and load balancing. Includes transmit and receive load balancing for IPV4 traffic. Receive load balancing is achieved through ARP negotiation. Requires Ethtool support in the base drivers and should be supported by all SimplStor network cards.

Next select the IPv4 addressing you plan to use. Automatic will get the network parameters from DHCP while Manual allows you to enter static IP settings.



If designating Manual you will need:

- IPv4 address and netmask prefix.
- Default gateway.
- IP address(es) of DNS servers.
- DNS domains to search.
- Any custom routes you require.

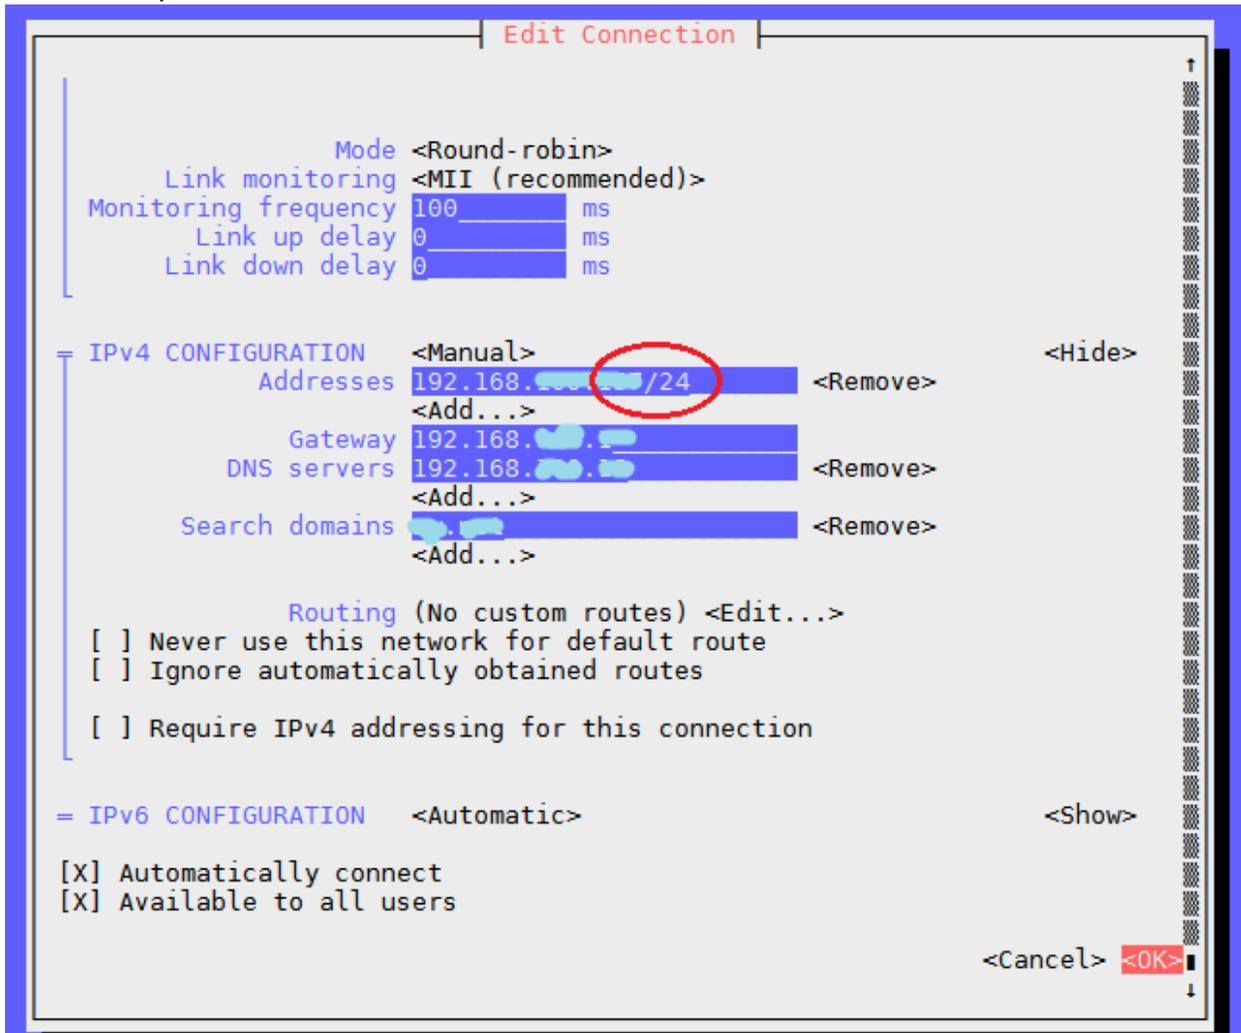
SimplStor Division:

sales@simplstor.com

support@simplstor.com

All rights reserved. The information contained herein is subject to change without notice. Reproduction, adaptation, or translation without prior written permission is prohibited, except as allowed under the copyright laws. SimplStor is a trademark of Zerowait Corporation. The Zerowait Logo, ZHA, and Zerowait are registered trademarks of the Zerowait Corporation. ©2017 Zerowait Corporation. Page 8

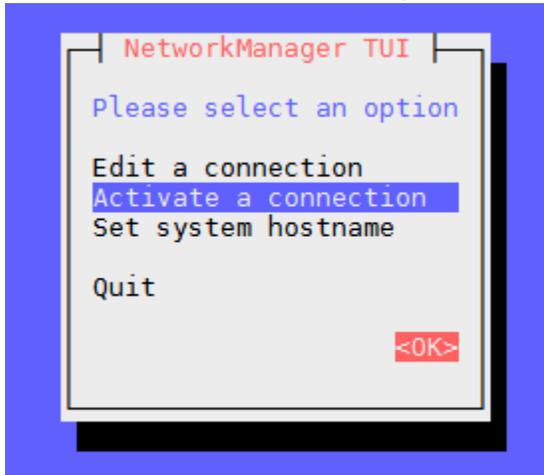
If using Manual, enter all the needed parameters. Be sure to include the network mask in CIDR notation as part of the IP address.



After completing these steps, saving and exiting “nmtui”, you can now examine the “ifcfg” files that were created in the /etc/sysconfig/network-scripts directory. If you need to make changes you can use the “nmtui” utility or you can make minor edits to the “ifcfg” files themselves.

Activate the connection and test

The simplest way to activate a connection is to use the “nmtui” utility. Select “Activate a connection” and on the subsequent screen select your newly created “bond0” interface.



Restart networking service should not be necessary, but if desired or needed use systemctl:

```
# systemctl restart network
```

Verify everything is working:

```
# cat /proc/net/bonding/bond0
Ethernet Channel Bonding Driver: v3.7.1 (April 27, 2011)

Bonding Mode: load balancing (round-robin)
MII Status: up
MII Polling Interval (ms): 100
Up Delay (ms): 0
Down Delay (ms): 0

Slave Interface: enp1s0f0
MII Status: up
Speed: 1000 Mbps
Duplex: full
Link Failure Count: 0
Permanent HW addr: 00:25:90:9b:xx:yy
Slave queue ID: 0

Slave Interface: enp1s0f1
MII Status: up
Speed: 1000 Mbps
Duplex: full
Link Failure Count: 0
Permanent HW addr: 00:25:90:9b:xx:yz
Slave queue ID: 0
```

SimplStor Division:

sales@simplstor.com

support@simplstor.com

All rights reserved. The information contained herein is subject to change without notice. Reproduction, adaptation, or translation without prior written permission is prohibited, except as allowed under the copyright laws. SimplStor is a trademark of Zerowait Corporation. The Zerowait Logo, ZHA, and Zerowait are registered trademarks of the Zerowait Corporation. ©2017 Zerowait Corporation. Page 10

Ping your default gateway to test your connectivity:

```
# ping 192.168.xxx.yyy
```

List all interfaces:

```
# ip link show
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN
mode DEFAULT qlen 1
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: enpls0f0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq
master bond0 state UP mode DEFAULT qlen 1000
    link/ether 00:25:90:9b:xx:zz brd ff:ff:ff:ff:ff:ff
3: enpls0f1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq
master bond0 state UP mode DEFAULT qlen 1000
    link/ether 00:25:90:9b:xx:zz brd ff:ff:ff:ff:ff:ff
4: bond0: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 1500 qdisc
noqueue state UP mode DEFAULT qlen 1000
    link/ether 00:25:90:9b:xx:zz brd ff:ff:ff:ff:ff:ff
5: virbr0: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc noqueue
state DOWN mode DEFAULT qlen 1000
    link/ether 52:54:00:4e:xy:az brd ff:ff:ff:ff:ff:ff
6: virbr0-nic: <BROADCAST,MULTICAST> mtu 1500 qdisc pfifo_fast master
virbr0 state DOWN mode DEFAULT qlen 1000
    link/ether 52:54:00:4e:xy:az brd ff:ff:ff:ff:ff:ff
```

The final test is to reboot your system and see that the interfaces come up as expected.

```
# systemctl reboot
```

At this point you should have Ethernet channel bonding working on your SimplStor!

Team Interfaces

Red Hat Enterprise Linux 7 also includes an alternate technology called “teaming” or “teams” that uses very lightweight kernel techniques. If you are interested in this please read the RHEL 7 Networking Guide for further details. A utility called “bond2team” can be used to convert your “bond” configuration to “teams” if you decide to explore this option. You can also use “nmtui” to create Team interfaces directly.

The following command will deliver the ifcfg files into a temporary directory where you can review the files before copying to the right location. To convert the current "bond0" ifcfg configuration to team ifcfg:

```
# /usr/bin/bond2team --master bond0
```

SimplStor Division:

sales@simplstor.com

support@simplstor.com

All rights reserved. The information contained herein is subject to change without notice. Reproduction, adaptation, or translation without prior written permission is prohibited, except as allowed under the copyright laws. SimplStor is a trademark of Zerowait Corporation. The Zerowait Logo, ZHA, and Zerowait are registered trademarks of the Zerowait Corporation. ©2017 Zerowait Corporation. Page 11