

[SLE] SuSE 9.0 Channel Bonding broken, fix and Mini-Howto

Source: <http://linux.derkeiler.com/Mailing-Lists/SuSE/2004-08/4068.html>

From: Randy Rue (rrue_at_fhcr.org)

Date: 08/30/04

Date: Mon, 30 Aug 2004 09:32:00 -0700

To: suse-linux-e@suse.com

I apologize for having misposted this on my first try. I've posted it again (hopefully correctly this time) so that it will be available to anyone searching on it. I've been searching on this topic myself for a few days, and have found a lot of other questions but not many answers.

– rrue

Hello All,

After testing it appears that Channel Bonding as shipped in SuSE 9.0 is broken. The NIC team appears to load normally but does not survive when the wrong cable is disconnected. Specifically, the team switches over without a hitch but whichever slave interface was added to the team last will not function, and the team won't function until the first NIC is reconnected AND the second NIC disconnected (forcing the team to pick the first NIC back up). This was confirmed by reversing the order of the bonded slave interfaces attached in the "ifenslave" command: changing the command from 'ifenslave bond0 eth0 eth1' to ifenslave bond0 eth1 eth0' made the problem reverse itself.

It looks like the default compile of ifenslave included with SuSE 9.0 ignores the instructions in `/usr/src/**kernel_version**/Documentation/networking/bonding.txt` and compiles using the header file `if_bonding.h` found in `/usr/include/linux` instead of the proper one in the kernel source directory `/usr/src/linux/include`. Recompiling and replacing the executable makes it work. A running ping session stumbles briefly when the active NIC is disconnected, but only a few packets are dropped. After about 20–30 seconds, the connection stabilizes (I have then let it run for as many as 1000 pings without dropping any packets).

The good news is that the proper version of ifenslave now also supports mode 6 (adaptive load balancing), providing both fault-tolerance and aggregated throughput.

SuSE: [SLE] SuSE 9.0 Channel Bonding broken, fix and Mini-Howto

I have attached a mini-howto below.

Randy Rue

Randall Rue
System Administrator, Postmaster
IT, Server Operations
Fred Hutchinson Cancer Research Center
Seattle, WA
USA

NETWORK INTERFACE TEAMING USING KERNEL-BASED CHANNEL BONDING IN SUSE 9.0

Version 2.0 August 27, 2004

1. Recompile the executable `/sbin/ifenslave` (SuSE's default compile of this command calls the wrong header files, and will appear to load normally but does not function properly):
 - a. Find the file `"ifenslave.c"` using `"locate ifenslave.c"` and `cd` to its directory.
 - b. `gcc -Wall -Wstrict-prototypes -O -I/usr/src/linux/include ifenslave.c -o ifenslave`
 - c. `mv /sbin/ifenslave /sbin/ifenslave.orig`
 - d. `mv ifenslave /sbin/ifenslave`
2. Add the following two lines to `/etc/modules.conf.local`:

```
alias bond0 bonding
options bonding miimon=100 mode=6
```
3. Use Yast to configure both Ethernet interfaces as DHCP clients.
4. Edit `/etc/sysconfig/network/ifcfg-ethX` for each of the interfaces `eth0` and `eth1` (or whatever physical interfaces are to be bonded):
 - a. Retain the `"Unique='xxxxyyzzz'"` defined by Yast
 - b. Otherwise modify the files as shown:

```
DEVICE='ethX'
USERCTL='no'
ONBOOT='yes'
STARTMODE='onboot'
MASTER='bond0'
SLAVE='yes'
BOOTPROTO='none'
UNIQUE='xxxxyyzzz'
```
5. Create a file `/etc/sysconfig/network/ifcfg-bond0` with the same permissions as the other `ifcfg-xxx` files:

```
DEVICE='bond0'
IPADDR='140.107.xxx.yyy'
NETMASK='255.255.25x.0'
```

SuSE: [SLE] SuSE 9.0 Channel Bonding broken, fix and Mini-Howto

```
NETWORK='140.107.xxx.0'  
BROADCAST='140.107.xxx.255'  
STARTMODE='onboot'  
BOOTPROTO='static'  
USERCTL='no'  
BONDING_MASTER='yes'  
BONDING_SLAVE0='eth0'  
BONDING_SLAVE1='eth1'
```

6. Edit /etc/init.d/network: to the end of the "start" case statement (approximately line 200), add the following two lines as shown.

Before:

```
esac  
rc_reset  
done  
[ ! -z "$WAIT_FOR_INTERFACES" ] && sleep $WAIT_FOR_INTERFACES
```

After:

```
esac  
rc_reset  
done  
## Added to start NIC bonding when starting network ##  
/sbin/ifenslave bond0 eth0 eth1  
[ ! -z "$WAIT_FOR_INTERFACES" ] && sleep $WAIT_FOR_INTERFACES
```

7. Restart the network interfaces ("rcnetwork restart") and test for proper configuration using ping, ifconfig and "cat /proc/net/bond0/info"

8. Reboot the machine and test again for proper configuration.

9. Start a ping (i.e. ping yahoo.com) and test that the process recovers and continues when either interface is disconnected.

--

Check the headers for your unsubscription address
For additional commands send e-mail to suse-linux-e-help@suse.com
Also check the archives at <http://lists.suse.com>
Please read the FAQs: suse-linux-e-faq@suse.com