Linux Bonding and VLANs With IBM BladeCenter

- RHEL 5 - SLES 10 & 11



Authors: William S. Champion wchamp@us.ibm.com

> Jason Daniel jdaniel@us.ibm.com

Nathan Flowers nflowers@us.ibm.com

Version: 1.0.0 Date: October 5, 2010

© IBM Corporation 2010

Revision History1.0.0 - October 5, 2010Initial Release

Notices:

This paper is intended to provide information regarding IBM[®] System x[™] Power Uninterruptable Power Supply (UPS). It discusses findings based on configurations that were created and tested under laboratory conditions. These findings may not be realized in all customer environments, and implementation in such environments may require additional steps, configurations and performance analysis. The information herein is provided "AS IS" with no warranties, express or implied. This information does not constitute a specification or form part of the warranty for any IBM or non-IBM products.

Information in this document was developed in conjunction with the use of the equipment specified and is limited in application to those specific hardware and software products and levels.

The information contained in this document has not been submitted to any formal IBM test and is distributed **as is**. The use of this information or the implementation of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

IBM may not officially support techniques mentioned in this document. For questions regarding officially supported techniques, please refer to the product documentation, announcement letters or contact the IBM Support Line at 1-800-IBM-SERV.

This document makes references to vendor-acquired applications or utilities. It is the customer responsibility to obtain licenses of these utilities prior to their usage.

© Copyright International Business Machines Corporation 2010. All rights reserved. U.S. Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Table of Contents

TABLE OF CONTENTS	
CONFIGURING LINUX BONDING & VLANS ON RHEL 5	4
BONDING CONFIGURATION VLAN CONFIGURATION IFCONFIG Depug	
CONFIGURING LINUX BONDING & VLANS ON SLES 10	
Bonding Configuration VLAN Configuration IFCONFIG Debug	
CONFIGURING LINUX BONDING & VLANS ON SLES 11	14
BONDING CONFIGURATION VLAN CONFIGURATION IFCONFIG DEBUG	
NIC BONDING PARAMETERS	19
MODE=	
CISCO CATALYST SWITCH MODULES (3110X) CONFIGURATION	21
BNT 1/10GB OPEINK SWITCH MODULES CONFIGURATION	
BNT STACKING FUNCTION CONFIGURATION BNT VLAN SWITCH CONFIGURATION	
STACK CABLING	26
Cisco 3110X Switch Modules Stack BNT 1/10Gb Switch Modules Stack	

Configuring Linux Bonding & VLANs on RHEL 5

Assume you want to configure VLAN20, VLAN 30, VLAN 40, VLAN 50, and VLAN 100 on a bonded interface that includes the eth0 and eth1 NICs on an HS22 blade. In this example VLAN 100 is untagged and the remaining VLANs are tagged. Further assume as an example that the desired static IP address assignment is as follows:

VLAN 20	tagged	192.168.20.5
VLAN 30	tagged	192.168.30.5
VLAN 40	tagged	192.168.40.5
VLAN 50	tagged	192.168.50.5
VLAN100	untagged	192.168.100.5

NOTE: Execute commands listed in this document as root.

Bonding Configuration

Add the lines in **blue** to the **modprobe**.conf file:

```
/etc/modprobe.conf
```

```
alias eth0 bnx2
alias eth1 bnx2
alias bond0 bonding
alias scsi_hostadapter mptbase
alias scsi_hostadapter1 mptsas
alias usb0 usbnet
```

Create the bond interface by creating the *ifcfg-bond0* file as below. To assign an IP address to the untagged VLAN, assign an IP address to the bond0 interface (*ifcfg-bond0*) and configure the switch modules with a Native VLAN. If no untagged VLAN is required then omit the IP address.

```
/etc/sysconfig/network-scripts/ifcfg-bond0
```

```
DEVICE=bond0
BONDING_OPTS="mode=4 miimon=500"
BOOTPROTO=none
ONBOOT=yes
NETWORK=192.168.100.0
NETMASK=255.255.255.0
IPADDR=192.168.100.5
USERCTL=no
TYPE=BOND
IPV6INIT=no
```

Note: The ONBOOT option for the bond interface should be set to yes so that this interface starts automatically.

See section "<u>NIC Bonding Parameters</u>" for information on BONDING OPTS parameters.

Bind eth0 to bond0 as a slave changing the ifcfg-eth0 file as follows:

```
/etc/sysconfig/network-scripts/ifcfg-eth0
    DEVICE=eth0
    BOOTPROTO=none
    ONBOOT=yes
    MASTER=bond0
    SLAVE=yes
    USERCTL=no
    TYPE=Ethernet
```

Bind eth1 to bond0 as a slave changing the ifcfg-eth1 file as follows:

```
/etc/sysconfig/network-scripts/ifcfg-eth1
```

```
DEVICE=eth1
BOOTPROTO=none
ONBOOT=yes
MASTER=bond0
SLAVE=yes
USERCTL=no
TYPE=Ethernet
```

Note: The ONBOOT option for the NIC interfaces can also be set to NO. If this is done, you may see a message during startup that this interface was skipped. When the bond interface activates, the physical interfaces should become active.

VLAN Configuration

Create the tagged VLAN 20 interface by creating the *ifcfg-bond0.20* file as follows:

/etc/sysconfig/network-scripts/ifcfg-bond0.20

DEVICE=bond0.20 BOOTPROTO=none ONBOOT=yes NETWORK=192.168.20.0 NETMASK=255.255.255.0 IPADDR=192.168.20.5 USERCTL=no VLAN=yes TYPE=Ethernet

Create the tagged VLAN 30 interface by creating the *ifcfg-bond0.30* file as follows:

/etc/sysconfig/network-scripts/ifcfg-bond0.30

```
DEVICE=bond0.30
BOOTPROTO=none
ONBOOT=yes
NETWORK=192.168.30.0
NETMASK=255.255.255.0
IPADDR=192.168.30.5
USERCTL=no
VLAN=yes
TYPE=Ethernet
```

Create the tagged VLAN 40 interface by creating the *ifcfg-bond0.40* file as follows:

```
/etc/sysconfig/network-scripts/ifcfg-bond0.40
    DEVICE=bond0.40
    BOOTPROTO=none
    ONBOOT=yes
    NETWORK=192.168.40.0
    NETMASK=255.255.255.0
    IPADDR=192.168.40.5
    USERCTL=no
    VLAN=yes
    TYPE=Ethernet
```

Create the tagged VLAN 50 interface by creating the ifcfg-bond0.50 file as follows:

```
/etc/sysconfig/network-scripts/ifcfg-bond0.50
```

```
DEVICE=bond0.50
BOOTPROTO=none
ONBOOT=yes
NETWORK=192.168.50.0
NETMASK=255.255.255.0
IPADDR=192.168.50.5
USERCTL=no
VLAN=yes
TYPE=Ethernet
```

Load the VLAN tagging module with modprobe command:

modprobe 8021q

Add the VLANs to the /proc/net/vlan filesystem with vconfig command:

vconfig add bond0 20 vconfig add bond0 30 vconfig add bond0 40 vconfig add bond0 50

Restart networking on the blade via the following command, and you are finished.

service network restart

IFCONFIG

The **ifconfig** command should now look similar to this:

- bond0 Link encap:Ethernet HWaddr 00:1A:64:AE:AE:A6
 inet addr:192.168.100.5 Bcast:192.168.100.255 Mask:255.255.255.0
 inet6 addr: fe80::21a:64ff:feae:aea6/64 Scope:Link
 UP BROADCAST RUNNING MASTER MULTICAST MTU:1500 Metric:1
 RX packets:55408 errors:0 dropped:0 overruns:0 frame:0
 TX packets:7109 errors:0 dropped:0 overruns:0 carrier:0
 collisions:0 txqueuelen:0
 RX bytes:3723363 (3.5 MiB) TX bytes:627309 (612.6 KiB)
- bond0.20 Link encap:Ethernet HWaddr 00:1A:64:AE:AE:A6
 inet addr:192.168.20.5 Bcast:192.168.20.255 Mask:255.255.255.0
 inet6 addr: fe80::21a:64ff:feae:aea6/64 Scope:Link
 UP BROADCAST RUNNING MASTER MULTICAST MTU:1500 Metric:1
 RX packets:266936 errors:0 dropped:0 overruns:0 frame:0
 TX packets:434 errors:0 dropped:0 overruns:0 carrier:0
 collisions:0 txqueuelen:0
 RX bytes:13348990 (12.7 MiB) TX bytes:56594 (55.2 KiB)
- bond0.30 Link encap:Ethernet HWaddr 00:1A:64:AE:AE:A6 inet addr:192.168.30.5 Bcast:192.168.30.255 Mask:255.255.255.0 inet6 addr: fe80::21a:64ff:feae:aea6/64 Scope:Link UP BROADCAST RUNNING MASTER MULTICAST MTU:1500 Metric:1 RX packets:265227 errors:0 dropped:0 overruns:0 frame:0 TX packets:178 errors:0 dropped:0 overruns:0 carrier:0 collisions:0 txqueuelen:0 RX bytes:13262218 (12.6 MiB) TX bytes:27639 (26.9 KiB)
- bond0.40 Link encap:Ethernet HWaddr 00:1A:64:AE:AE:A6
 inet addr:192.168.40.5 Bcast:192.168.40.255 Mask:255.255.255.0
 inet6 addr: fe80::21a:64ff:feae:aea6/64 Scope:Link
 UP BROADCAST RUNNING MASTER MULTICAST MTU:1500 Metric:1
 RX packets:265216 errors:0 dropped:0 overruns:0 frame:0
 TX packets:178 errors:0 dropped:0 overruns:0 carrier:0
 collisions:0 txqueuelen:0
 RX bytes:13261750 (12.6 MiB) TX bytes:27625 (26.9 KiB)
- bond0.50 Link encap:Ethernet HWaddr 00:1A:64:AE:AE:A6
 inet addr:192.168.50.5 Bcast:192.168.50.255 Mask:255.255.255.0
 inet6 addr: fe80::21a:64ff:feae:aea6/64 Scope:Link
 UP BROADCAST RUNNING MASTER MULTICAST MTU:1500 Metric:1
 RX packets:265212 errors:0 dropped:0 overruns:0 frame:0
 TX packets:173 errors:0 dropped:0 overruns:0 carrier:0
 collisions:0 txqueuelen:0
 RX bytes:13261482 (12.6 MiB) TX bytes:27150 (26.5 KiB)
- eth0 Link encap:Ethernet HWaddr 00:1A:64:AE:AE:A6 UP BROADCAST RUNNING SLAVE MULTICAST MTU:1500 Metric:1 RX packets:41868 errors:0 dropped:0 overruns:0 frame:0 TX packets:6555 errors:0 dropped:0 overruns:0 carrier:0 collisions:0 txqueuelen:1000 RX bytes:2784247 (2.6 MiB) TX bytes:556589 (543.5 KiB) Interrupt:193 Memory:fa000000-fa012800

```
eth1 Link encap:Ethernet HWaddr 00:1A:64:AE:AE:A6
UP BROADCAST RUNNING SLAVE MULTICAST MTU:1500 Metric:1
RX packets:13540 errors:0 dropped:0 overruns:0 frame:0
TX packets:554 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:939116 (917.1 KiB) TX bytes:70720 (69.0 KiB)
Interrupt:130 Memory:92000000-92012800
lo Link encap:Local Loopback
inet addr:127.0.0.1 Mask:255.0.0.0
inet6 addr: ::1/128 Scope:Host
UP LOOPBACK RUNNING MTU:16436 Metric:1
RX packets:1446 errors:0 dropped:0 overruns:0 frame:0
```

```
TX packets:1446 errors:0 dropped:0 overruns:0 riame.0
TX packets:1446 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:2235018 (2.1 MiB) TX bytes:2235018 (2.1 MiB)
```

Debug

1. To view added VLANs run the following command:

```
cat /proc/net/vlan/config
```

```
VLAN Dev name| VLAN IDName-Type:VLAN_NAME_TYPE_RAW_PLUS_VID_NO_PADbond0.20| 20bond0.30| 30bond0.40| 40bond050bond099bond0(Unwanted VLAN)
```

If you created a VLAN with the wrong VLAN ID, remove it from the /proc/net/vlan filesystem with vconfig. (Restarting the network alone will not suffice to remove the bogus VLAN).

```
vconfig rem bond0.99
```

Removed VLAN -: bond0.99:-

2. Via "Ismod", ensure the following modules are running. For example:

8021q	57425	0
bonding	138201	0

3. Do not forget to add the VLANs into the configuration for whatever Internal Switch Modules you are using in the BladeCenter. Be sure you carry all the VLANs on your server facing ports.

Configuring Linux Bonding & VLANs on SLES 10

Assume you want to configure VLAN20, VLAN 30, VLAN 40, VLAN 50, and VLAN 100 on a bonded interface that includes the eth0 and eth1 NICs on an HS22 blade. In this example VLAN 100 is untagged and the remaining VLANs are tagged. Further assume as an example that the desired static IP address assignment is as follows:

VLAN 20	tagged	192.168.20.5
VLAN 30	tagged	192.168.30.5
VLAN 40	tagged	192.168.40.5
VLAN 50	tagged	192.168.50.5
VLAN100	untagged	192.168.100.5

NOTE: Execute commands listed in this document as root.

Bonding Configuration

SLES 10 uses the MAC Address as part of the interface configuration filename, ifcfg-eth-idxx:xx:xx:xx:xx:xx (the x's represent the MAC Address of the respective adapters) and is unique for each NIC.

For example, eth0 will look similar to this:

```
/etc/sysconfig/network/ifcfg-eth-id-xx:xx:xx:xx:xx
```

```
BOOTPROTO='none'
STARTMODE='auto'
NAME='NetXtreme II BCM5709S Gigabit Ethernet'
BROADCAST=''
ETHTOOL_OPTIONS=''
IPADDR=''
MTU=''
NETMASK=''
NETWORK=''
REMOTE_IPADDR=''
USERCONTROL='no'
```

For example, eth1 will look similar to this:

```
/etc/sysconfig/network/ifcfg-eth-id-xx:xx:xx:xx:xx:xx:
BOOTPROTO='none'
STARTMODE='auto'
NAME='NetXtreme II BCM5709S Gigabit Ethernet'
BROADCAST=''
ETHTOOL_OPTIONS=''
IPADDR=''
MTU=''
NETMASK=''
NETWORK=''
REMOTE_IPADDR=''
USERCONTROL='no'
```

Note: The STARTMODE option for the NIC interfaces can also be set to `off'. If this is done, you may see a message during startup that this interface was skipped. When the bond interface activates, the physical interfaces should become active.

Create the bond interface by creating the *ifcfg-bond0* file as below. To assign an IP address to the untagged VLAN, assign an IP address to the bond0 interface (*ifcfg-bond0*) and configure the switch modules with a Native VLAN. If no untagged VLAN is required then omit the IP address.

/etc/sysconfig/network/ifcfg-bond0

```
BONDING_MASTER='yes'
BONDING_MODULE_OPTS='mode=802.3ad miimon=500'
BONDING_SLAVE0='eth-id-xx:xx:xx:xx:xx'
BONDING_SLAVE1='eth-id-xx:xx:xx:xx:xx'
BOOTPROTO='none'
BROADCAST=''
BROADCAST=''
ETHTOOL_OPTIONS=''
IPADDR='192.168.100.5'
MTU=''
NAME=''
NETMASK='255.255.255.0'
NETWORK='192.168.100.0'
REMOTE_IPADDR=''
STARTMODE='auto'
USERCONTROL='no'
```

Note: The STARTMODE option for bond interface should be set to `auto' so that this interface starts automatically.

See section "<u>NIC Bonding Parameters</u>" for information on BONDING MODULE OPTS parameters.

VLAN Configuration

Create the tagged VLAN 20 interface by creating the ifcfg-vlan20 file as follows:

/etc/sysconfig/network/ifcfg-vlan20

```
BOOTPROTO='static'
BROADCAST=''
ETHERDEVICE='bond0'
ETHTOOL_OPTIONS=''
IPADDR='192.168.20.5/24'
MTU=''
NAME=''
NETWORK=''
REMOTE_IPADDR=''
STARTMODE='auto'
USERCONTROL='no'
```

Create the tagged VLAN 30 interface by creating the ifcfg-vlan30 file as follows:

/etc/sysconfig/network/ifcfg-vlan30

```
BOOTPROTO='static'
BROADCAST=''
ETHERDEVICE='bond0'
ETHTOOL_OPTIONS=''
IPADDR='192.168.30.5/24'
MTU=''
NAME=''
NETWORK=''
REMOTE_IPADDR=''
STARTMODE='auto'
USERCONTROL='no'
```

Create the tagged VLAN 40 interface by creating the ifcfg-vlan40 file as follows:

/etc/sysconfig/network/ifcfg-vlan40

```
BOOTPROTO='static'
BROADCAST=''
ETHERDEVICE='bond0'
ETHTOOL_OPTIONS=''
IPADDR='192.168.40.5/24'
MTU=''
NAME=''
NETWORK=''
REMOTE_IPADDR=''
STARTMODE='auto'
USERCONTROL='no'
```

Create the tagged VLAN 50 interface by creating the *ifcfg-vlan50* file as follows:

/etc/sysconfig/network/ifcfg-vlan50

```
BOOTPROTO='static'
BROADCAST=''
ETHERDEVICE='bond0'
ETHTOOL_OPTIONS=''
IPADDR='192.168.50.5/24'
MTU=''
NAME=''
NAME=''
NETWORK=''
REMOTE_IPADDR=''
STARTMODE='auto'
USERCONTROL='no'
```

The STARTMODE option for the Vlan interfaces should be set to `auto' so that these interfaces start automatically.

Restart networking on the blade via the following command, and you are finished.

```
service network restart
```

Additional options may be found in the etc/sysconfig/network/ifcfg.template file.

IFCONFIG

The **ifconfig** command should now look similar to this:

bond0	Link encap:Ethernet HWaddr 00:14:5E:D6:18:5E inet addr:192.168.100.5 Bcast:192.168.100.255 Mask:255.255.255.0 UP BROADCAST RUNNING MASTER MULTICAST MTU:1500 Metric:1 RX packets:55849 errors:0 dropped:54817 overruns:0 frame:0 TX packets:6144 errors:0 dropped:0 overruns:0 carrier:0 collisions:0 txqueuelen:0 RX bytes:3693326 (3.5 Mb) TX bytes:497803 (486.1 Kb)
eth0	Link encap:Ethernet HWaddr 00:14:5E:D6:18:5E UP BROADCAST RUNNING SLAVE MULTICAST MTU:1500 Metric:1 RX packets:42082 errors:0 dropped:41082 overruns:0 frame:0 TX packets:5567 errors:0 dropped:0 overruns:0 carrier:0 collisions:0 txqueuelen:1000 RX bytes:2746054 (2.6 Mb) TX bytes:424587 (414.6 Kb) Interrupt:82 Memory:da000000-da012100
ethl	Link encap:Ethernet HWaddr 00:14:5E:D6:18:5E UP BROADCAST RUNNING SLAVE MULTICAST MTU:1500 Metric:1 RX packets:13767 errors:0 dropped:13735 overruns:0 frame:0 TX packets:577 errors:0 dropped:0 overruns:0 carrier:0 collisions:0 txqueuelen:1000 RX bytes:947272 (925.0 Kb) TX bytes:73216 (71.5 Kb) Interrupt:90 Memory:d800000-d8012100
lo	Link encap:Local Loopback inet addr:127.0.0.1 Mask:255.0.0.0 UP LOOPBACK RUNNING MTU:16436 Metric:1 RX packets:155 errors:0 dropped:0 overruns:0 frame:0 TX packets:155 errors:0 dropped:0 overruns:0 carrier:0 collisions:0 txqueuelen:0 RX bytes:11608 (11.3 Kb) TX bytes:11608 (11.3 Kb)
vlan20	Link encap:Ethernet HWaddr 00:14:5E:D6:18:5E inet addr:192.168.20.5 Bcast:192.168.20.255 Mask:255.255.255.0 UP BROADCAST RUNNING MASTER MULTICAST MTU:1500 Metric:1 RX packets:21 errors:0 dropped:0 overruns:0 frame:0 TX packets:60 errors:0 dropped:0 overruns:0 carrier:0 collisions:0 txqueuelen:0 RX bytes:1302 (1.2 Kb) TX bytes:3080 (3.0 Kb)
vlan30	Link encap:Ethernet HWaddr 00:14:5E:D6:18:5E inet addr:192.168.30.5 Bcast:192.168.30.255 Mask:255.255.255.0 UP BROADCAST RUNNING MASTER MULTICAST MTU:1500 Metric:1 RX packets:21 errors:0 dropped:0 overruns:0 frame:0 TX packets:18 errors:0 dropped:0 overruns:0 carrier:0 collisions:0 txqueuelen:0 RX bytes:1344 (1.3 Kb) TX bytes:1372 (1.3 Kb)
vlan40	Link encap:Ethernet HWaddr 00:14:5E:D6:18:5E inet addr:192.168.40.5 Bcast:192.168.40.255 Mask:255.255.255.0 UP BROADCAST RUNNING MASTER MULTICAST MTU:1500 Metric:1 RX packets:19 errors:0 dropped:0 overruns:0 frame:0 TX packets:16 errors:0 dropped:0 overruns:0 carrier:0 collisions:0 txqueuelen:0

RX bytes:1176 (1.1 Kb) TX bytes:1176 (1.1 Kb)

vlan50 Link encap:Ethernet HWaddr 00:14:5E:D6:18:5E inet addr:192.168.50.5 Bcast:192.168.50.255 Mask:255.255.255.0 UP BROADCAST RUNNING MASTER MULTICAST MTU:1500 Metric:1 RX packets:21 errors:0 dropped:0 overruns:0 frame:0 TX packets:18 errors:0 dropped:0 overruns:0 carrier:0 collisions:0 txqueuelen:0 RX bytes:1302 (1.2 Kb) TX bytes:1316 (1.2 Kb)

Debug

1. To view added VLANs run the following command:

```
cat /proc/net/vlan/config
```

VLAN Dev name | VLAN ID Name-Type: VLAN_NAME_TYPE_RAW_PLUS_VID_NO_PAD bond0.20 | 20 | bond0 bond0.30 | 30 | bond0 bond0.40 | 40 | bond0 bond0.50 | 50 | bond0

2. Via "Ismod", ensure the following modules are running. For example:

8021q	57425	0
bonding	138201	0

3. Do not forget to add the VLANs into the configuration for whatever Internal Switch Modules you are using in the BladeCenter. Be sure you carry all the VLANs on your server facing ports.

Configuring Linux Bonding & VLANs on SLES 11

Assume you want to configure VLAN20, VLAN 30, VLAN 40, VLAN 50, and VLAN 100 on a bonded interface that includes the eth0 and eth1 NICs on an HS22 blade. In this example VLAN 100 is untagged and the remaining VLANs are tagged. Further assume as an example that the desired static IP address assignment is as follows:

VLAN 20	tagged	192.168.20.5
VLAN 30	tagged	192.168.30.5
VLAN 40	tagged	192.168.40.5
VLAN 50	tagged	192.168.50.5
VLAN100	untagged	192.168.100.5

NOTE: Execute commands listed in this document as root.

Bonding Configuration

For example, eth0 will look similar to this:

```
/etc/sysconfig/network/ifcfg-eth0
```

```
BOOTPROTO='none'
STARTMODE='auto'
NAME='NetXtreme II BCM5709S Gigabit Ethernet'
BROADCAST=''
ETHTOOL_OPTIONS=''
IPADDR=''
MTU=''
NETMASK=''
NETWORK=''
REMOTE_IPADDR=''
USERCONTROL='no'
```

For example, eth1 will look similar to this:

/etc/sysconfig/network/ifcfg-eth1

```
BOOTPROTO='none'
STARTMODE='auto'
NAME='NetXtreme II BCM5709S Gigabit Ethernet'
BROADCAST=''
ETHTOOL_OPTIONS=''
IPADDR=''
MTU=''
NETMASK=''
NETWORK=''
REMOTE_IPADDR=''
USERCONTROL='no'
```

Note: The STARTMODE option for the NIC interfaces can also be set to `off'. If this is done, you may see a message during startup that this interface was skipped. When the bond interface activates, the physical interfaces should become active.

Create the bond interface by creating the *ifcfg-bond0* file as below. To assign an IP address to the untagged VLAN, assign an IP address to the bond0 interface (*ifcfg-bond0*) and configure the switch modules with a Native VLAN. If no untagged VLAN is required then omit the IP address.

/etc/sysconfig/network/ifcfg-bond0

```
BONDING_MASTER='yes'
BONDING_MODULE_OPTS='mode=802.3ad miimon=500'
BONDING_SLAVE0='eth0'
BONDING_SLAVE1='eth1'
BOOTPROTO='static'
BROADCAST='192.168.100.255'
ETHTOOL_OPTIONS=''
IPADDR='192.168.100.5'
MTU=''
NAME=''
NETMASK='255.255.255.0'
NETWORK='192.168.100.0'
REMOTE_IPADDR=''
STARTMODE='auto'
USERCONTROL='no'
```

Note: The STARTMODE option for bond interface should be set to `auto' so that this interface starts automatically.

See section "<u>NIC Bonding Parameters</u>" for information on BONDING MODULE OPTS parameters.

VLAN Configuration

Create the tagged VLAN 20 interface by creating the ifcfg-vlan20 file as follows:

```
/etc/sysconfig/network/ifcfg-vlan20
```

```
BOOTPROTO='static'
BROADCAST=''
ETHERDEVICE='bond0'
ETHTOOL_OPTIONS=''
IPADDR='192.168.20.5/24'
MTU=''
NAME=''
NETWORK=''
REMOTE_IPADDR=''
STARTMODE='auto'
USERCONTROL='no'
```

Create the tagged VLAN 30 interface by creating the *ifcfg-vlan30* file as follows:

```
/etc/sysconfig/network/ifcfg-vlan30
```

```
BOOTPROTO='static'
BROADCAST=''
ETHERDEVICE='bond0'
ETHTOOL_OPTIONS=''
IPADDR='192.168.30.5/24'
MTU=''
NAME=''
NAME=''
NETWORK=''
REMOTE_IPADDR=''
STARTMODE='auto'
USERCONTROL='no'
```

Create the tagged VLAN 40 interface by creating the ifcfg-vlan40 file as follows:

```
/etc/sysconfig/network/ifcfg-vlan40
```

```
BOOTPROTO='static'
BROADCAST=''
ETHERDEVICE='bond0'
ETHTOOL_OPTIONS=''
IPADDR='192.168.40.5/24'
MTU=''
NAME=''
NETWORK=''
REMOTE_IPADDR=''
STARTMODE='auto'
USERCONTROL='no'
```

Create the tagged VLAN 50 interface by creating the *ifcfg-vlan50* file as follows:

```
/etc/sysconfig/network/ifcfg-vlan50
```

```
BOOTPROTO='static'
BROADCAST=''
ETHERDEVICE='bond0'
ETHTOOL_OPTIONS=''
IPADDR='192.168.50.5/24'
MTU=''
NAME=''
NETWORK=''
REMOTE_IPADDR=''
STARTMODE='auto'
USERCONTROL='no'
```

The STARTMODE option for the Vlan interfaces should be set to `auto' so that these interfaces start automatically.

Restart networking on the blade via the following command, and you are finished.

```
service network restart
```

Additional options may be found in the etc/sysconfig/network/ifcfg.template file.

IFCONFIG

The **ifconfig** command should now look similar to this:

bond0	Link encap:Ethernet HWaddr 00:14:5E:D5:E1:56 inet addr:192.168.100.5 Bcast:192.168.100.255 Mask:255.255.255.0 UP BROADCAST RUNNING MASTER MULTICAST MTU:1500 Metric:1 RX packets:53984 errors:0 dropped:0 overruns:0 frame:0 TX packets:5244 errors:0 dropped:0 overruns:0 carrier:0 collisions:0 txqueuelen:0 RX bytes:3564893 (3.3 Mb) TX bytes:420101 (410.2 Kb)
eth0	Link encap:Ethernet HWaddr 00:14:5E:D5:E1:56 UP BROADCAST RUNNING SLAVE MULTICAST MTU:1500 Metric:1 RX packets:40119 errors:0 dropped:0 overruns:0 frame:0 TX packets:4665 errors:0 dropped:0 overruns:0 carrier:0 collisions:0 txqueuelen:1000 RX bytes:2610429 (2.4 Mb) TX bytes:346309 (338.1 Kb) Interrupt:17 Memory:da000000-da012800
ethl	Link encap:Ethernet HWaddr 00:14:5E:D5:E1:56 UP BROADCAST RUNNING SLAVE MULTICAST MTU:1500 Metric:1 RX packets:13865 errors:0 dropped:0 overruns:0 frame:0 TX packets:579 errors:0 dropped:0 overruns:0 carrier:0 collisions:0 txqueuelen:1000 RX bytes:954464 (932.0 Kb) TX bytes:73792 (72.0 Kb) Interrupt:19 Memory:d8000000-d8012800
lo	Link encap:Local Loopback inet addr:127.0.0.1 Mask:255.0.0.0 UP LOOPBACK RUNNING MTU:16436 Metric:1 RX packets:56 errors:0 dropped:0 overruns:0 frame:0 TX packets:56 errors:0 dropped:0 overruns:0 carrier:0 collisions:0 txqueuelen:0 RX bytes:5216 (5.0 Kb) TX bytes:5216 (5.0 Kb)
vlan20	Link encap:Ethernet HWaddr 00:14:5E:D5:E1:56 inet addr:192.168.20.5 Bcast:192.168.20.255 Mask:255.255.255.0 UP BROADCAST RUNNING MASTER MULTICAST MTU:1500 Metric:1 RX packets:22 errors:0 dropped:0 overruns:0 frame:0 TX packets:19 errors:0 dropped:0 overruns:0 carrier:0 collisions:0 txqueuelen:0 RX bytes:1386 (1.3 Kb) TX bytes:1414 (1.3 Kb)
vlan30	Link encap:Ethernet HWaddr 00:14:5E:D5:E1:56 inet addr:192.168.30.5 Bcast:192.168.30.255 Mask:255.255.255.0 UP BROADCAST RUNNING MASTER MULTICAST MTU:1500 Metric:1 RX packets:21 errors:0 dropped:0 overruns:0 frame:0 TX packets:18 errors:0 dropped:0 overruns:0 carrier:0 collisions:0 txqueuelen:0 RX bytes:1344 (1.3 Kb) TX bytes:1372 (1.3 Kb)
vlan40	Link encap:Ethernet HWaddr 00:14:5E:D5:E1:56 inet addr:192.168.40.5 Bcast:192.168.40.255 Mask:255.255.255.0 UP BROADCAST RUNNING MASTER MULTICAST MTU:1500 Metric:1 RX packets:21 errors:0 dropped:0 overruns:0 frame:0 TX packets:18 errors:0 dropped:0 overruns:0 carrier:0 collisions:0 txqueuelen:0

RX bytes:1344 (1.3 Kb) TX bytes:1372 (1.3 Kb)

vlan50 Link encap:Ethernet HWaddr 00:14:5E:D5:E1:56 inet addr:192.168.50.5 Bcast:192.168.50.255 Mask:255.255.255.0 UP BROADCAST RUNNING MASTER MULTICAST MTU:1500 Metric:1 RX packets:20 errors:0 dropped:0 overruns:0 frame:0 TX packets:17 errors:0 dropped:0 overruns:0 carrier:0 collisions:0 txqueuelen:0 RX bytes:1302 (1.2 Kb) TX bytes:1330 (1.2 Kb)

Debug

1. To view added VLANs run the following command:

```
cat /proc/net/vlan/config
```

VLAN Dev name | VLAN ID Name-Type: VLAN_NAME_TYPE_RAW_PLUS_VID_NO_PAD bond0.20 | 20 | bond0 bond0.30 | 30 | bond0 bond0.40 | 40 | bond0 bond0.50 | 50 | bond0

2. Via "Ismod", ensure the following modules are running. For example:

8021q	57425	0
bonding	138201	0

3. Do not forget to add the VLANs into the configuration for whatever Internal Switch Modules you are using in the BladeCenter. Be sure you carry all the VLANs on your server facing ports.

NIC Bonding Parameters

mode=

The configuration instructions in this document uses 4 or ' 802.3ad' for the mode parameter. All the mode parameter options are documented below.

Round-robin policy: Packets are transmitted in sequential order from the first available slave through the last.

RHEL: mode=0 SLES: mode='balance-rr'

Features: load balancing and fault tolerance

Active-backup policy: Only one bonded slave is active. Another slave becomes active only if the current active slave fails. Because this is an active-backup configuration, only one interface is active one the network at any time.

RHEL: mode=1 SLES: mode=' active-backup'

Features: fault tolerance

XOR policy: The interface matches the inbound MAC address with the MAC address for one of the slave NICs. When the link is established, packets are transmitted in sequence, beginning with the first available interface.

RHEL: mode=2 SLES: mode=' balance-xor'

Features: load balancing and fault tolerance

Broadcast policy: All packets are transmitted on all slave interfaces. This mode provides fault tolerance.

RHEL: mode=3
SLES: mode=' broadcast'

Features: fault tolerance

IEEE 802.3ad Dynamic link aggregation: This mode is also known as LACP. An aggregation group is created that shares the same speed and duplex settings. In accordance with 802.3ad specification all slaves are utilized in the active bond.

RHEL: mode=4 SLES: mode=' 802.3ad'

Features: load balancing and fault tolerance

Pre-requisites:

- 1. Support in the base NIC drivers for retrieving the speed and duplex of each slave via ethtool.
- 2. A switch that supports IEEE 802.3ad Dynamic link aggregation. The switch, the NICs are connected to, will need to be configured to enable 802.3ad mode.

Adaptive transmit load balancing: No special switch support is required to support this mode. The out bound packets are distributed according to the current load on each slave. In bound packets are received by the current slave. Another slave takes over if the receiving slave fails.

RHEL: mode=6 SLES: mode=' balance-tlb'

Features: load balancing and fault tolerance

Prerequisite:

Support in the base NIC drivers for retrieving the speed and duplex of each slave via ethtool.

Adaptive load balancing: No special switch support is required to support this mode. Includes balance-tlb plus receive load balancing (rlb) for IPV4 traffic. The inbound load balancing is achieved by ARP negotiation.

RHEL: mode=6 SLES: mode=' balance-alb'

Features: load balancing and fault tolerance

miimon=

The milmon parameter used in the configuration instructions is 500. Depending on the network environment, the milmon parameter may need to be adjusted.

miimon= Specifies (in milliseconds) how often MII link monitoring transpires. This is useful if high availability is required because MII is used to verify that the NIC is active.

To verify that the base NIC driver supports the MII tool, type the following command:

ethtool <interface-name> | grep "Link detected:"

<interface-name> will need to be replaced with the name of the NIC interface, such as eth0, in the command above. If MII is supported, the command returns:

Link detected: yes

Cisco Catalyst Switch Modules (3110X) Configuration

The following is a sample configuration for using the IBM BladeCenter with Cisco Catalyst switch modules (3110X). The configuration steps assume four blades servers are in server slots 4, 5, 6 & 7 and the Ethernet switch modules are in the factory default configuration and have been installed in I/O Bays 1 and 2 of the BladeCenter Chassis. The switch modules are interconnected via the stack cable to form a single logical switching unit that spans the two switches. See <u>Cisco 3110X</u> Switch Modules Stack section for details on stack cabling. Below are the command line interface steps, to configure the Cisco switch modules.

Cisco VLAN Switch Configuration

To configure the VLANs on the switch:

```
Switch>enable
Password:
Switch#configure terminal
```

Enter configuration commands, one per line. End with CNTL/Z.

```
Switch(config) #vlan 10
Switch(config-vlan) #name Management
Switch(config-vlan) #vlan 20
Switch(config-vlan) #name Private
Switch(config-vlan) #vlan 30
Switch(config-vlan) #vlan 40
Switch(config-vlan) #vlan 40
Switch(config-vlan) #vlan 50
Switch(config-vlan) #vlan 50
Switch(config-vlan) #name DMZ
Switch(config-vlan) #name DMZ
Switch(config-vlan) #vlan 100
Switch(config-vlan) #name ServerMgmt
Switch(config-vlan) #name ServerMgmt
```



Note regarding VLANs:

In this configuration, VLAN 10 is the management VLAN for AMM and switch management. This
VLAN should not include any servers in order to separate management traffic from production traffic.
VLAN 100 has been configured as a direct server management VLAN to maintain isolation of AMM
traffic from server traffic.

The two NICs of each blade will be configured using NIC bonding in LACP mode. The switch ports for each pair of NICs must also be configured to use a Port Channel in LACP mode.

```
Switch(config)#interface gigabitEthernet 1/0/4
Switch(config-if)# channel-group 4 mode active
Switch(config-if)# channel-group 4 mode active
Switch(config-if)# channel-group 4 mode active
Switch(config-if)#exit
Switch(config)#interface Port-channel4
Switch(config-if)#switchport
Switch(config-if)#switchport trunk allowed vlan 20,30,40,50,100
Switch(config-if)#switchport trunk native vlan 100
Switch(config-if)# switchport mode trunk
Switch(config-if)# switchport mode trunk
Switch(config-if)# exit
Switch(config-if)#exit
Switch(config-if)# channel-group 5 mode active
Switch(config-if)#interface gigabitEthernet 2/0/5
```

```
Switch(config-if) # channel-group 5 mode active
Switch(config-if) #exit
Switch(config)#interface Port-channel5
Switch(config-if)#switchport
Switch(config-if)#switchport trunk allowed vlan 20,30,40,50,100
Switch(config-if)#switchport trunk native vlan 100
Switch(config-if) # switchport mode trunk
Switch (config) #interface gigabitEthernet 1/0/6
Switch(config-if) # channel-group 6 mode active
Switch(config-if) #interface gigabitEthernet 2/0/6
Switch(config-if) # channel-group 6 mode active
Switch(config-if)#exit
Switch (config) #interface Port-channel6
Switch(config-if)#switchport
Switch(config-if)#switchport trunk allowed vlan 20,30,40,50,100
Switch(config-if) #switchport trunk native vlan 100
Switch(config-if) # switchport mode trunk
Switch(config-if)#exit
Switch (config) #interface gigabitEthernet 1/0/7
Switch(config-if)# channel-group 7 mode active
Switch(config-if) #interface gigabitEthernet 2/0/7
Switch(config-if) # channel-group 7 mode active
Switch(config-if)#exit
Switch(config)#interface Port-channel7
Switch(config-if)#switchport
Switch(config-if)#switchport trunk allowed vlan 20,30,40,50,100
Switch(config-if)#switchport trunk native vlan 100
Switch(config-if) # switchport mode trunk
```

Switch(config-if)#exit
Switch(config)#port-channel load-balance src-dst-ip

Notes regarding interface configuration:

- In this example, NIC bonding for each server is configured to transmit tagged frames for VLANs 20, 30, 40 and 50. The bonding is also configured to forward untagged frames from the NICs with the native VLAN configured as VLAN 100 on the port channel interface and VLAN 100 added to the list of allowed VLANs.
- Configuration of interfaces providing connectivity to upstream switches has intentionally been omitted from this example as the requirements are dependent upon the configuration of the upstream switch(es). The interfaces should be configured in a manner that would be compatible operational mode and allow traffic of all of the necessary VLANs to traverse the link.
- The default port channel load balancing is defaults to use source MAC address to balance the traffic on the link. In a route network environment, this is typically not an optimal setting. In this example, the setting has been changed to source / destination IP.

BNT 1/10Gb Uplink Switch Modules Configuration

The following is a sample configuration for using the IBM BladeCenter with BNT 1/10Gb Uplink Ethernet switch modules. The configuration steps assume four blades servers are in server slots 4, 5, 6 & 7 and the Ethernet switch modules are in the factory default configuration and have been installed in I/O Bays 1 and 2 of the BladeCenter Chassis. The switch modules are interconnected via a fiber cable to form a single logical switching unit that spans the two switches using the staking function of the BNT modules. See <u>BNT 1/10Gb</u> <u>Switch Modules Stack</u> section for details on stack cabling. Below are the command line interface steps, to configure the BNT switch modules.

BNT Stacking Function Configuration

To configure the stacking function:

Note: The UUID included in these commands is the UUID of the BladeCenter chassis in which the switch modules are installed.

BNT VLAN Switch Configuration

To configure the VLANs on the switch:

```
>> Main#
/c/12/vlan 10
        enable
        name "Management"
        def 0
/c/12/vlan 20
        enable
        name "Private"
        def 1:4-1:7 2:4-2:7
/c/12/vlan 30
        enable
        name "Signaling"
        def 1:4-1:7 2:4-2:7
/c/12/vlan 40
        enable
        name "RTP"
        def 1:4-1:7 2:4-2:7
/c/12/vlan 50
        enable
        name "DMZ"
        def 1:4-1:7 2:4-2:7
/c/12/vlan 100
        enable
        name "ServerMqmt"
        def 1:4-1:7 2:4-2:7
/c/l2/stg 1/clear
/c/l2/stg 1/add 1 10 20 30 40 50 100 4090
```



Note regarding VLANs:

- In this configuration, VLAN 10 is the management VLAN for AMM and switch management. This
 VLAN should not include any servers in order to separate management traffic from production traffic.
 VLAN 100 has been configured as a direct server management VLAN to maintain isolation of AMM
 traffic from server traffic.
- In this example VLAN 10 shows not external uplinks; in live network the VLAN would include the uplink interfaces to upstream switches.
- VLAN 4090 is listed in the spanning tree commands but is not listed in the VLAN configuration.
 VLAN 4090 is the default stacking VLAN that is created when the BNT modules are configured in stacking mode.

The two NICs of each blade will be configured using NIC bonding in LACP mode. The switch ports for each pair of NICs must also be configured to use a Port Channel in LACP mode.

>> Main# /c/port 1:4 pvid 100 /c/port 1:5 pvid 100 /c/port 1:6 pvid 100 /c/port 1:7 pvid 100 /c/port 2:4 pvid 100 /c/port 2:5 pvid 100 /c/port 2:6 pvid 100 /c/port 2:7 pvid 100 /c/12/lacp/port 1:4 mode active adminkey 4 /c/l2/lacp/port 1:5 mode active adminkey 5 /c/l2/lacp/port 1:6 mode active adminkey 6 /c/12/lacp/port 1:7 mode active adminkey 7 /c/12/lacp/port 2:4 mode active adminkev 4 /c/l2/lacp/port 2:5 mode active adminkey 5 /c/l2/lacp/port 2:6 mode active adminkey 6 /c/12/lacp/port 2:7 mode active adminkey 7

Notes regarding interface configuration:

- In this example, NIC bonding for each server is configured to transmit tagged frames for VLANs 20, 30, 40 and 50. The bonding is also configured to forward untagged frames from the NICs with the PVID configured 100 on the interfaces for the blade servers.
- Configuration of interfaces providing connectivity to upstream switches has intentionally been omitted from this example as the requirements are dependent upon the configuration of the upstream switch(es). The interfaces should be configured in a manner that would be compatible operational mode and allow traffic of all of the necessary VLANs to traverse the link.
- The default port channel load balancing is defaults to use source MAC address to balance the traffic on the link. In a route network environment, this is typically not an optimal setting. In this example, the setting has been changed to source / destination IP.

Stack Cabling

Cisco 3110X Switch Modules Stack

See <u>Cisco Catalyst Switch Modules (3110X) Configuration</u> section for details on configuring the Cisco switch modules for LACP and Port Channel.

BCH w/ Cisco 3110X

BCE w/ Cisco 3110X



BNT 1/10Gb Switch Modules Stack

See <u>BNT 1/10Gb Uplink Switch Modules Configuration</u> section for details on configuring the Cisco switch modules for LACP and Port Channel.



BCH w/ BNT 1/10Gb

BCE w/ BNT 1/10Gb



IBM, the IBM Logo, BladeCenter, and System x are registered trademarks or trademarks of International Business Machines Corporation in the United States, other countries or both.

Other company, product or service names may be trademarks or service marks of others.