

# Assigning Interrupts to Processor Cores using an Intel® 82575 or 82598 Ethernet Controller

---

*February 2009*



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, life sustaining, critical control or safety systems, or in nuclear facility applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Intel Corporation may have patents or pending patent applications, trademarks, copyrights, or other intellectual property rights that relate to the presented subject matter. The furnishing of documents and other materials and information does not provide any license, express or implied, by estoppel or otherwise, to any such patents, trademarks, copyrights, or other intellectual property rights.

**IMPORTANT - PLEASE READ BEFORE INSTALLING OR USING INTEL® PRE-RELEASE PRODUCTS.**

Please review the terms at [http://www.intel.com/netcomms/prerelease\\_terms.htm](http://www.intel.com/netcomms/prerelease_terms.htm) carefully before using any Intel® pre-release product, including any evaluation, development or reference hardware and/or software product (collectively, "Pre-Release Product"). By using the Pre-Release Product, you indicate your acceptance of these terms, which constitute the agreement (the "Agreement") between you and Intel Corporation ("Intel"). In the event that you do not agree with any of these terms and conditions, do not use or install the Pre-Release Product and promptly return it unused to Intel.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See [http://www.intel.com/products/processor\\_number](http://www.intel.com/products/processor_number) for details.

The 82575 and 82598 Ethernet Controllers may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Hyper-Threading Technology requires a computer system with an Intel® Pentium® 4 processor supporting HT Technology and a HT Technology enabled chipset, BIOS and operating system. Performance will vary depending on the specific hardware and software you use. See [http://www.intel.com/products/ht/Hyperthreading\\_more.htm](http://www.intel.com/products/ht/Hyperthreading_more.htm) for additional information.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an ordering number and are referenced in this document, or other Intel literature, may be obtained from:


Intel Corporation  
P.O. Box 5937  
Denver, CO 80217-9808

or call in North America 1-800-548-4725, Europe 44-0-1793-431-155, France 44-0-1793-421-777, Germany 44-0-1793-421-333, other Countries 708-296-9333.

Intel and Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

\*Other names and brands may be claimed as the property of others.

Copyright © 2009, Intel Corporation. All Rights Reserved.



## Contents

---

- 1.0 IP Forwarding and Pinning Queues to Processors** ..... 5
  - 1.1 Assumptions ..... 5
  - 1.2 Build a New Kernel ..... 5
  - 1.3 Install the New Kernel and Restart the System ..... 7
  - 1.4 IGB and IXGBE Drivers ..... 7
    - 1.4.1 IGB ..... 7
    - 1.4.2 IXGBE ..... 9
  - 1.5 Initializing IP Forwarding ..... 11
- 2.0 Other Optimizations Found From Performance Testing** ..... 11


## Revision History

---

Rev	Rev Date	Description
2.0	Feb 2009	Initial public release.



*Note: This page intentionally left blank.*



## 1.0 IP Forwarding and Pinning Queues to Processors

---

Multicore processors and the newest Ethernet adapters (including the 82575 and 82598) allow TCP forwarding flows to be optimized by assigning execution flows to individual cores. By default, Linux automatically assigns interrupts to processor cores. Two methods currently exist for automatically assigning the interrupts, an in-kernel IRQ balancer and the IRQ balance daemon in user space. Both offer tradeoffs that might lower CPU usage but do not maximize the IP forwarding rates. Optimal throughput can be obtained by manually pinning the queues of the Ethernet adapter to specific processor cores.

For IP forwarding, a transmit/receive queue pair should use the same processor core and reduce any cache synchronization between different cores. This can be performed by assigning transmit and receive interrupts to specific cores. With Linux kernels 2.6.23 and greater, multiple queues can be used on the 82575 and the 82598. Also, Extended Messaging Signaled Interrupts (MSI-X) support a larger number of interrupts that can be used, allowing for finer-grained control and targeting of the interrupts to specific CPUs.

*Note:* The cores on one socket of a dual-socket system can be assigned non-consecutive numbers. For example, on a dual-socket quad-core system Socket 0 has CPU 0, 2, 4, 6 and Socket 1 has CPU 1, 3, 5, 7 on x86\_64, and Socket 0 has CPU 0, 1, 2, 3 and Socket 1 has CPU 4, 5, 6, 7 on x86\_32.

### 1.1 Assumptions

The following hardware and software are assumed in this document's IP forwarding example:

- x86\_64 SMP system
- Intel Ethernet adapter using 82575EB or 82598EB Ethernet controllers
- Installed Fedora Core 8 x86\_64 with development packages
- Kernel version 2.6.7.7-134.fc10.x86\_64 (at least kernel version 2.6.23 for MSI-X)
- Network connection

### 1.2 Build a New Kernel

*Note:* These instructions were summarized from

<http://fedoraproject.org/wiki/Docs/CustomKernel>.

Build the kernel as a user with "sudo" privileges.

1. Install prerequisite packages *rpmdevtools* and *yum-utils*.

```
sudo yum install yum-utils rpmdevtools
```



2. Prepare an RPM package building environment in a home directory.

```
rpmdev-setuptree
```

3. Download the kernel source rpm file.

```
yumdownloader --source kernel
```

4. Install build dependencies for the kernel source:

```
sudo yum-builddep kernel-2.6.27.7-134.fc10.src.rpm'
```

5. Install the kernel into the local *rpmbuild* directories. Note that kernel-2.6.27.7-134.fc10.src.rpm installs into *~/rpmbuild/SOURCES* and *~/rpmbuild/SPECS*.

```
rpm -Uvh kernel-2.6.27.7-134.fc10.src.rpm
```

Ignore any messages similar to the following:

```
warning: user mockbuild does not exist - using root
```

```
warning: group mockbuild does not exist - using root
```

6. Prepare the kernel source tree. The source is in the *~/rpmbuild/BUILD/kernel-2.6.24/linux-2.6.24.x86\_64*.

```
cd ~/rpmbuild/SPECS
```

```
rpmbuild -bp --target=`uname -m` kernel.spec
```

```
cd ~/rpmbuild/BUILD/kernel-2.6.247/linux-2.6.247.x86_64/
```

7. Copy the default configuration file.

```
cp configs/kernel-2.6.27-7-x86_64.config .config
```

```
make oldconfig
```

8. Configure the kernel.

```
make menuconfig
```

9. Lower the interrupt frequency. 100 Hz is a typical choice for servers, SMP and NUMA with most processors that might show reduced performance because too many timer interrupts are occurring.

```
Processor type and features ---> Timer Frequency ---> 100Hz
```

10. Turn off kernel packet filtering.


```
Networking Support ---> Networking options ---> Network packet filtering  
framework (Netfilter) ---> N
```

11. Add a new comment line to the top of the *.config* file.

The comment should contain the hardware platform the kernel is built for. This is found with *uname -i*.

For example, an *x86\_64* machine has the following line added to the top of the config file:

```
# x86_64
```



12. Copy the generated config file to `~/rpmbuild/SOURCES/`:

```
cp .config ~/rpmbuild/SOURCES/config-x86_64
```

13. Edit the `kernel.spec` file to create unique kernel package names.

```
cd ~/rpmbuild/SPECS
```

14. Edit the `buildid` line in the `kernel.spec` file to generate a unique kernel name. Change the line `## define buildid .local` to `%define buildid .test`.

Note that there is a space between the `"%"` and `"define"` that needs to be deleted.

15. Use the `rpmbuild` utility to build the new RPM files.

```
rpmbuild -bb --target=`uname -m` kernel.spec
```

The build process takes some time to complete. If the build completes successfully, the new kernel packages can be found in the `~/rpmbuild/RPMS` directory.

## 1.3 Install the New Kernel and Restart the System

1. Install the new kernel with the `rpm -ivh` command. Do not use the `-U` or `--upgrade` options as it might be necessary to go to an older version of the kernel if problems occur.

```
rpm -ivh ~/rpmbuild/RPMS/x86_64/kernel- 2.6.27.7-134.test.fc10.x86_64.rpm
```

2. Reboot.

```
shutdown -r now
```

## 1.4 IGB and IXGBE Drivers

The following sections are specific to the `igb` or the `ixgbe` drivers.

### 1.4.1 IGB

1. Get the `igb` driver from sourceforge, compile and install.

Note that the extra CFLAGS allocate a separate handler for Tx cleanups. (This assumes an Ethernet adapter that uses the `igb` driver.)

```
wget http://downloads.sourceforge.net/e1000/igb-1.2.44.9.tar.gz
```

```
tar xzf igb-1.2.44.9.tar.gz
```

```
cd igb-1.2.44.9/src/
```

```
sudo make CFLAGS_EXTRA="-DCONFIG_IGB_SEPARATE_TX_HANDLER " install
```

```
cd ../../
```

2. Remove the old drivers.


```
rmmod igb
```

3. Add the module to the kernel using MSI-X interrupts and multiple queues.

```
modprobe igb IntMode=3
```

4. Make sure IRQ balance is turned off. This fails if the service isn't running, but the error can be ignored.

```
service irqbalance stop
```



5. Make sure iptables is off. This fails if the service isn't running, but the error can be ignored.

```
service iptables stop
```

6. Make sure the interface is up. The interface must be up for the interrupts to appear in */proc/interrupts*.

Check *ifconfig -all* to see which interfaces are assigned to the Ethernet adapter. This example assumes that the interfaces are *eth1* and *eth2*.

```
ifconfig eth1 up
```

```
ifconfig eth2 up
```

7. Assign interrupts to specific processors. Check the number of processor (cores) by checking */proc/cpuinfo*.

```
cat /proc/cpuinfo | grep processor
```

The output looks like:

```
processor      : 0
processor      : 1
processor      : 2
processor      : 3
```

8. Check the *proc filesystem* for interrupts.

The following table should appear in */proc/interrupts*:

2292:	3	0	0	0	PCI-MSI-edge	eth2
2293:	78	0	0	0	PCI-MSI-edge	eth2-rx0
2294:	3	0	0	0	PCI-MSI-edge	eth2-tx0
2295:	17	0	0	0	PCI-MSI-edge	eth1
2296:	89	0	0	0	PCI-MSI-edge	eth1-rx0
2297:	105	0	0	0	PCI-MSI-edge	eth1-tx0
2298:	1838	0	0	0	PCI-MSI-edge	eth0

The interrupts can be assigned to specific processors by assigning a processor to the interrupts. The processor number is converted to a binary mask, so processor 0 is 1, processor 1 is 2, processor 2 is 4, processor 3 is 8 ( $2^{\text{processor\_number}}$ ).


To assign *eth1-tx0* (2297) to processor 1, echo the converted mask to */proc/irq/2297/smp\_affinity*.

```
echo 2 > /proc/irq/2297/smp_affinity
```

**Note:**

The cores on one socket of a dual-socket system are assigned non-consecutive numbers. For example, on a dual-socket quad-core system Socket 0 has CPU 0, 2, 4, 6 and Socket 1 has CPU 1, 3, 5, 7.





## 1.4.2 IXGBE

1. Get the Ethernet driver from sourceforge, compile and install.

The extra CFLAGS disable Large Receive Offload (LRO). LRO increases performance by effectively letting the kernel handle fewer packet headers for the same amount of data but it causes the stack to fragment packets in the IP forwarding/routing case due to the large (effective) MTU size of the received packets.

Force compilation of NAPI (Rx polling mode).

*Note:* The shell variable `EXTRA_INC_DIR_SAVE` should be set to the `ioatdma` source directory when the `ioatdma` driver was installed.

```
wget http://downloads.sourceforge.net/e1000/ixgbe-1.3.31.5.tar.gz
tar xzf ixgbe-1.3.31.5.tar.gz
cd ixgbe-1.3.31.5/src
sudo make "CFLAGS_EXTRA=-DIXGBE_NO_LRO -DIXGBE_NAPI" install
cd ../../
```

2. Remove the old drivers.

```
rmmod ixgbe
```

3. Add the module to the kernel. If available on the platform, add the module to the kernel using Receive Side Scaling (RSS). `RSS=1` sets the descriptor queue count to the number of online processors up to a maximum of 16.

Use packet split mode for Rx to receive the packet header in the header buffer and the payload in the data buffer by setting `RxBufferMode` to 1.

```
modprobe ixgbe RSS=1 RxBufferMode=1
```

4. Make sure `IRQ balance` is off. This fails if the service isn't running, but the error can be ignored.

```
service irqbalance stop
```

5. Make sure `iptables` is off. This fails if the service isn't running, but the error can be ignored.

```
service iptables stop
```

6. Make sure the interface is up. The interface must be up for the interrupts to appear in `/proc/interrupts`. This example assumes that the interfaces are `eth1` and `eth2`.

```
ifconfig eth1 up
```

```
ifconfig eth2 up
```

7. Reduce the transmit ring to 256 entries and the receive ring to 512 entries.

```
ethtool -G tx 256
```

```
ethtool -G rx 512
```

8. Assign interrupts to specific processors. Check the number of processor (cores) by checking `/proc/cpuinfo`.

```
cat /proc/cpuinfo | grep processor
```



The output looks like:

```
processor      : 0
processor      : 1
processor      : 2
processor      : 3
```

9. Check the *proc filesystem* for interrupts.

The following table should appear in */proc/interrupts*:

2280:	0	0	0	0	PCI-MSI-edge	eth2:lsc
2281:	6	0	0	0	PCI-MSI-edge	eth2:v4-Tx
2282:	6	0	0	0	PCI-MSI-edge	eth2:v3-Rx
2283:	6	0	0	0	PCI-MSI-edge	eth2:v2-Rx
2284:	6	0	0	0	PCI-MSI-edge	eth2:v1-Rx
2285:	6	0	0	0	PCI-MSI-edge	eth2:v0-Rx
2286:	0	0	0	0	PCI-MSI-edge	eth1:lsc
2287:	7	0	0	0	PCI-MSI-edge	eth1:v4-Tx
2288:	7	0	0	0	PCI-MSI-edge	eth1:v3-Rx
2289:	7	0	0	0	PCI-MSI-edge	eth1:v2-Rx
2290:	7	0	0	0	PCI-MSI-edge	eth1:v1-Rx
2291:	7	0	0	0	PCI-MSI-edge	eth3:v0-Rx


The interrupts can be assigned to specific processors by assigning a processor to the interrupts. The processor number is converted to a binary mask, so processor 0 is 1, processor 1 is 2, processor 2 is 4, processor 3 is 8 ( $2^{\text{processor\_number}}$ ).

To assign *eth3:v0-Rx (2290)* to processor 1, echo the converted mask to */proc/irq/2290/smp\_affinity*.

```
echo 2 > /proc/irq/2290/smp_affinity
```

*Note:*

The cores on one socket of a dual-socket system are assigned non-consecutive numbers. For example, on a dual-socket quad-core system Socket 0 has CPU 0, 2, 4, 6 and Socket 1 has CPU 1, 3, 5, 7.



## 1.5 Initializing IP Forwarding

IP forwarding is initialized by default on reboot.

1. Edit `/etc/sysctl.conf` and change the line:

```
net.ipv4.ip_forward = 0
```

to

```
net.ipv4.ip_forward = 1
```

2. Initialize IP forwarding immediately.

```
sudo sysctl -w net.ipv4.ip_forward=1
```

## 2.0 Other Optimizations Found From Performance Testing

Testing has shown increased performance with some stack tuning on Linux operating systems. The tuning mainly involves increasing resources for the network stack. This includes increasing the TCP window size.

These parameters are configurable using the `sysctl` utility in Linux as indicated in this example:

```
sysctl -w net.core.rmem_default_524287
```

Running `sysctl` without the `-w` argument lists the parameter with its current setting. Though increasing the values can help improve performance, it is necessary to experiment with different values to determine what works best for a given system, workload and traffic type.

Note that these are only suggested starting points.

```
/proc/sys/net/core/rmem_default - default receive window
```

```
(default=124928), suggested change to 524287
```

```
/proc/sys/net/core/wmem_default - default send window
```

```
(default=124928), suggested change to 524287
```

```
/proc/sys/net/core/rmem_max - maximum receive window
```

```
(default=131071), suggested change to 524287
```

```
/proc/sys/net/core/wmem_max - maximum send window
```

```
(default=131071), suggested change to 524287
```

```
/proc/sys/net/core/optmem_max - maximum option memory buffers
```

```
(default=20480), suggested change to 524287
```

```
/proc/sys/net/core/netdev_max_backlog - number of unprocessed input packets  
before kernel starts dropping them
```

```
(default=1000), suggested change to 300000
```

```
/proc/sys/net/ipv4/tcp_rmem - memory reserved for TCP rcv buffers (min default  
max)
```



```
(defaults 4096 87380 4194304), suggested change to 10000000 10000000  
10000000
```

```
/proc/sys/net/ipv4/tcp_wmem - memory reserved for TCP snd buffers (min default  
max)
```

```
(defaults 4096 16384 4194304), suggested change to 10000000 10000000  
10000000
```

```
/proc/sys/net/ipv4/tcp_mem - memory reserved for TCP buffers (min default max)
```

```
(defaults 193152 257536 386304), suggested change to 10000000 10000000  
10000000
```