

Linux Enumeration of NICs

Version 1.0

By Robert Hentosh (*robert_hentosh@dell.com*) (August 2006)

Linux enumeration of the network interfaces may not align with BIOS and chassis labeling of the Ethernet ports. The administrator can rename the Ethernet interfaces to coincide with the labeling of the ports.

Linux may not enumerate the network interfaces as expected. System administrators may expect the onboard network ports labeled LOM1, LOM2, etc... on a system to be assigned interface names eth0, eth1, etc... respectively. This would be one possible naming convention for network interfaces; however no industry standards currently exist to ensure such a convention. The method used in Linux to determine the number associated with a network port (i.e. eth0, eth1, eth2...) is complex and changed with the 2.6 and newer kernels.

How Linux assigns network interface names

The default name for Ethernet interfaces is based upon how Linux initializes them during device discovery. As Linux finds the network devices it will start numbering them starting with 0 and increasing sequentially. Device discovery is dependent on the device driver load order, PCI bus topology and the device driver code.

First, the device driver load order is determined by the */etc/modprobe.conf* file. Device drivers assigned lower interface numbers in that file are loaded first. When Linux loads a single device driver it will initialize and find all devices supported by that single driver first.

Next, the PCI bus topology is composed of buses, bridges and devices. PCI devices must be connected to a PCI bus. The PCI buses are connected by PCI bridges to either other buses or to the system. The topology of the system can be viewed as a tree. Using this analogy, the devices are leaves, the system is the trunk, buses are branches and bridges exist where branches meet each other or the trunk. Searching for PCI devices in a system is accomplished by "walking the tree". The method of walking the tree was modified in the 2.6 and later kernels, thus changing the order in which devices are found.

And last, each device driver will search the PCI tree for all the devices it supports. Some drivers have a list of different chips it will support and search the tree for each device in the list. Other drivers will scan the tree and, for each device, see if it is in its list of supported devices. This will also change order of how devices are found and thus its interface name.

Changes in system configuration will also result in a different enumeration order. If a new network card is inserted into the PCI tree, its new position could be between two previous network devices. This may result in the new card taking over the name of a previous card in the system.

There is currently no industry standard to allow the OS to determine the physical labeling of Ethernet ports on the motherboard.

Method of working around this issue

A system administrator can use the Ethernet hardware address (HWADDR) of a network device to assign a specific network interface name to a network port. Each Ethernet device has a unique hardware address associated with it. Hardware addresses are normally hex numbers written in the form "XX:XX:XX:XX:XX:XX".

Red Hat Enterprise Linux 3 and 4 contain configuration files */etc/sysconfig/network-scripts/ifcfg-`<network interface name>`*. These files contain a line, *HWADDR=*, which specifies

the hardware address of each port. This line can be modified, or added, if missing, to assign a specific network interface name to a specific network port. Read `/usr/share/doc/initscripts-*/sysconfig.txt` for more information.

SuSE Linux Enterprise Server 10 uses the configuration files in `/etc/udev/rules.d/` subdirectory to associate a hardware address with a Linux Ethernet device name. A configuration line in the file, `net_persistent_names.rules`, can be modified, or added, if missing, to assign a specific network interface name to a network hardware address associated to a specific network port. Below is an example line:

```
SUBSYSTEM=="net", ACTION=="add", SYSFS{address}=="00:02:b3:5b:75:f3",
IMPORT="/lib/udev/rename_netiface %k eth0"
```

Read `/usr/share/doc/packages/sysconfig/README.Persistent_Interface_Names` for more information.

An administrator can prevent network devices from being renamed when new Ethernet cards are added to a system by using the above methods. It will also allow the administrator to reconfigure the Linux device names to coincide with the names physically labeled on the machine.

Identifying the network interface's port

The `ethtool` command can be used to discover which port on a server is currently associated with which Linux network device name. Run the command:

```
ethtool -p eth0
```

to cause the port for the network interface device name given (e.g. `eth0`) to be identified in some manner. This typically results in the blinking of one or more of the LEDs associated with that Ethernet port.

Script to assist with network device naming

A script has been developed at Dell to assist with naming the network interfaces. It will attempt to rename the Ethernet ports in a more expected manner. You can find the script at http://linux.dell.com/files/name_eths/.