

Linux Enumeration of NICs

Version2.0

By Robert Hentosh <robert_hentosh@dell.com> and Matt Domsch <Matt_Domsch@dell.com>
October 2006

Abstract

Linux naming of the hardware network interfaces may not align with BIOS and chassis labeling of the Ethernet ports. This is seen on Dell PowerEdge 1950, 1955, 2900, and 2950 servers when using a Linux 2.6 kernel-based product such as Red Hat Enterprise Linux 4, and Novell/SuSE Linux Enterprise Server 9 and 10. Solutions to align the Linux name with the expected name are presented.

Introduction

System administrators may expect the onboard network ports labeled Gb1, Gb2, 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.

The naming disconnect can be observed at two different times:

1. Install Time - OS Installation when network booting (e.g. using a PXE installation procedure);
2. System Runtime – Following OS installation

Dell engineers have developed workarounds for Red Hat Enterprise Linux 4 (RHEL4), and Novell/SuSE Linux Enterprise Server (SLES) 9 and 10 for each of the above scenarios.

While the naming disconnect directly affects Dell PowerEdge 1950, 1955, 2900, and 2950 servers running Linux 2.6 kernels (eg. RHEL4, RHEL5 beta, SLES9, SLES10), this behavior is not unique to Dell systems. The workarounds and solutions presented here work on Dell systems, and may work on additional systems.

The naming disconnect issues are not observed on Dell servers when using 2.4 kernel based distributions such as RHEL2.1, RHEL3, Red Hat Linux 7.2 and 7.3, or SLES8.

Background

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.

On RHEL4, 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.

On SLES, device drivers are loaded based on the PCI bus topology as discovered by the kernel. 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 devices 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 a PCI slot, 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.

The root cause of NIC enumeration mismatches is that there is currently no industry standard to enable the OS to determine the physical labeling of Ethernet ports on the motherboard.

Workarounds

Workarounds have been identified for both install time and runtime NIC enumeration issues.

Install Time

At install time, the network interface used to download the OS kernel and initial ramdisk from the PXE server may not be the *default* interface (i.e., eth0) used by the OS installer to download the rest of the OS image. This leads to the installer trying to use the "wrong" interface to finish the download, often leading to a failed installation.

RHEL and SLES installers have long had the ability to be told *which* interface to use to download the rest of the image. On systems where the installer's default interface isn't the "right" interface, you can use these kernel command line options to resolve it:

- For RHEL4 Update 3 and higher, and in RHEL5 beta, when using a PXE server, the NIC used for the initial PXE download can be used for the rest of the OS installation as well. In your PXE server's configuration files, use these options:

```
IPAPPEND 2
APPEND ksdevice=bootif
```

The OS installer then uses the same interface, regardless of its name, to finish the download.

- RHEL 4 Gold through Update 2, pass on the kernel command line:

```
ksdevice=eth1
```

This causes the installer to use eth1 instead of the default eth0. RHEL has the additional option:

```
ksdevice=link
```

This causes the installer to use the first network device it finds that has "link" – i.e. is plugged into a network switch.

- For SLES 9 or 10, pass on the kernel command line:

```
netdevice=eth1
```

This causes the installer to use eth1 instead of the default eth0.

Of course, instead of the above, one could disable or remove all but one NIC in the system, such that there is only one NIC which the kernel will name eth0.

System Runtime

name_eths

Dell developed a program, **name_eths**, which reduces the complexity of the solutions presented here by automatically implementing the specific mechanisms outlined for each OS version. We strongly recommend you use **name_eths** to accomplish these tasks. You can find the program at

http://linux.dell.com/files/name_eths/.

Name_eths must be run once, and run again if network cards are added to or removed from the system.

Name_eths presently can configure network device names on RHEL3, 4, and 5 (beta), Fedora Core 6 and earlier, SLES9, and SLES10. For other Linux distributions, use one of the manual workarounds below.

Manual workarounds

A system administrator can use the Ethernet hardware address (MAC address) 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. MAC addresses are normally hexadecimal 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. The **name_eths** program will do this for you.

SLES9 uses the configuration file in */etc/sysconfig/network/ifcfg-eth-id-<mac address>* for each Ethernet port. Each file may contain a *PERSISTENT_NAME=* line which specifies the name to be given to the device. One difficulty with this is that, for SLES9 releases up to and including Service Pack 3, one cannot assign names to swap "eth0" and "eth1". Instead, you must rename both devices into another name space, e.g. "eth-external" and "eth-internal". By modifying a few of the startup scripts, this restriction can be lifted and you can use "eth0" and "eth1" as expected. The **name_eths** program provides these updated scripts and will modify the config files for you.

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, *30-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"
```

The **name_eths** program will do this for you.

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 physical NIC port on a server (either embedded NIC or PCI-e/PCI-X NIC adapter cards) is currently associated with which Linux network device name. Run the command:

```
ethtool -p eth0
```

to cause the port for the indicated network interface device (in this example, `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.

Long-Term Industry Standard Solution

At present there is no formal method by which BIOS may communicate to the operating system the names of the devices as it knows them. This leads to the above outlined naming disconnects and provided workarounds. Dell, through our participation in standards bodies such as the DMTF SMBIOS Working Group the ACPI Forum, is working towards a formal documented solution to this issue. Once adopted into the standard and implemented in both BIOS and operating system code, it is expected that a uniform and widely accepted solution will be provided to the industry.