

[PATCH 19-rc1] Fix typos in /Documentation : 'U-Z'

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-10/msg02222.html>

- *From:* Matt LaPlante <kernel1@xxxxxxxxxxxxxxxxxxx>
 - *Date:* Fri, 6 Oct 2006 09:50:31 -0400
-

This patch fixes typos in various Documentation txts. The patch addresses some words starting with the letters 'U-Z'.

Looks like I made it through the alphabet...just in time to start over again too! Maybe I can fit more profound fixes into the next round...? Time will tell. :)

Signed-off-by: Matt LaPlante <kernel1@xxxxxxxxxxxxxxxxxxx>

--

```
diff -ru a/Documentation/accounting/taskstats.txt b/Documentation/accounting/taskstats.txt
--- a/Documentation/accounting/taskstats.txt 2006-10-05 22:18:50.000000000 -0400
+++ b/Documentation/accounting/taskstats.txt 2006-10-05 23:00:07.000000000 -0400
@@ -122,12 +122,12 @@
```

However, maintaining per-process, in addition to per-task stats, within the kernel has space and time overheads. To address this, the taskstats code

- accumulates each exiting task's statistics into a process-wide data structure.
- When the last task of a process exits, the process level data accumulated also
- +accumulates each exiting task's statistics into a process-wide data structure.
- +When the last task of a process exits, the process level data accumulated also gets sent to userspace (along with the per-task data).

When a user queries to get per-tgid data, the sum of all other live threads in

- the group is added up and added to the accumulated total for previously exited
- +the group is added up and added to the accumulated total for previously exited threads of the same thread group.

Extending taskstats

```
diff -ru a/Documentation/block/biodoc.txt b/Documentation/block/biodoc.txt
--- a/Documentation/block/biodoc.txt 2006-10-05 22:18:50.000000000 -0400
+++ b/Documentation/block/biodoc.txt 2006-10-05 23:04:59.000000000 -0400
@@ -391,7 +391,7 @@
```

on to the generic block layer, only to be merged by the i/o scheduler when the underlying device was capable of handling the i/o in one shot. Also, using the buffer head as an i/o structure for i/os that didn't originate

- from the buffer cache unnecessarily added to the weight of the descriptors
- +from the buffer cache unnecessarily added to the weight of the descriptors which were generated for each such chunk.

The following were some of the goals and expectations considered in the @@ -403,14 +403,14 @@

for raw i/o.

ii. Ability to represent high-memory buffers (which do not have a virtual address mapping in kernel address space).

-iii.Ability to represent large i/os w/o unnecessarily breaking them up (i.e

+iii.Ability to represent large i/os w/o unnecessarily breaking them up (i.e greater than PAGE_SIZE chunks in one shot)

iv. At the same time, ability to retain independent identity of i/os from different sources or i/o units requiring individual completion (e.g. for latency reasons)

v. Ability to represent an i/o involving multiple physical memory segments (including non-page aligned page fragments, as specified via readv/writev) - without unnecessarily breaking it up, if the underlying device is capable of + without unnecessarily breaking it up, if the underlying device is capable of handling it.

vi. Preferably should be based on a memory descriptor structure that can be passed around different types of subsystems or layers, maybe even

diff -ru a/Documentation/DMA-API.txt b/Documentation/DMA-API.txt

--- a/Documentation/DMA-API.txt 2006-10-05 22:18:50.000000000 -0400

+++ b/Documentation/DMA-API.txt 2006-10-05 23:07:35.000000000 -0400

@@ -489,7 +489,7 @@

flags can be or'd together and are

DMA_MEMORY_MAP - request that the memory returned from

-dma_alloc_coherent() be directly writeable.

+dma_alloc_coherent() be directly writable.

DMA_MEMORY_IO - request that the memory returned from

dma_alloc_coherent() be addressable using read/write/memcpy_toio etc.

diff -ru a/Documentation/dvb/ci.txt b/Documentation/dvb/ci.txt

--- a/Documentation/dvb/ci.txt 2006-10-05 22:18:48.000000000 -0400

+++ b/Documentation/dvb/ci.txt 2006-10-05 22:56:38.000000000 -0400

@@ -71,7 +71,7 @@

The disadvantage is that the driver/hardware has to manage the rest. For the application programmer it would be as simple as sending/receiving an array to/from the CI ioctls as defined in the Linux DVB API. No changes -have been made in the API to accomodate this feature.

+have been made in the API to accommodate this feature.

* Why the need for another CI interface ?

@@ -102,7 +102,7 @@

implemented by most applications. Hence this area is revisited.

This CI interface is quite different in the case that it tries to

-accomodate all other CI based devices, that fall into the other categories

+accommodate all other CI based devices, that fall into the other categories.

This means that this CI interface handles the EN50221 style tags in the

[PATCH 19-rc1] Fix typos in /Documentation : 'U-Z'

Application layer only and no session management is taken care of by the

```
diff -ru a/Documentation/eisa.txt b/Documentation/eisa.txt
--- a/Documentation/eisa.txt 2006-10-05 22:18:52.000000000 -0400
+++ b/Documentation/eisa.txt 2006-10-05 22:38:22.000000000 -0400
@@ -62,7 +62,7 @@
```

bus_base_addr : slot 0 address on this bus

slots : max slot number to probe

force_probe : Probe even when slot 0 is empty (no EISA mainboard)

-dma_mask : Default DMA mask. Usually the bridge device dma_mask.

+dma_mask : Default DMA mask. Usually the bridge device dma_mask.

bus_nr : unique bus id, set by eisa_root_register

** Driver :

```
diff -ru a/Documentation/filesystems/adfs.txt b/Documentation/filesystems/adfs.txt
--- a/Documentation/filesystems/adfs.txt 2006-10-05 22:18:51.000000000 -0400
+++ b/Documentation/filesystems/adfs.txt 2006-10-05 22:50:02.000000000 -0400
@@ -3,7 +3,7 @@
```

uid=nnn All files in the partition will be owned by user id nnn. Default 0 (root).

- gid=nnn All files in the partition will be in group

+ gid=nnn All files in the partition will be in group

nnn. Default 0 (root).

ownmask=nnn The permission mask for ADFS 'owner' permissions will be nnn. Default 0700.

```
diff -ru a/Documentation/filesystems/configfs/configfs.txt b/Documentation/filesystems/configfs/configfs.txt
--- a/Documentation/filesystems/configfs/configfs.txt 2006-10-05 22:18:51.000000000 -0400
+++ b/Documentation/filesystems/configfs/configfs.txt 2006-10-05 22:42:08.000000000 -0400
@@ -209,7 +209,7 @@
```

[struct config_group]

-A config_item cannot live in a vacuum. The only way one can be created

+A config_item cannot live in a vacuum. The only way one can be created is via mkdir(2) on a config_group. This will trigger creation of a child item.

```
@@ -275,7 +275,7 @@
```

[struct configfs_subsystem]

-A subsystem must register itself, usually at module_init time. This

+A subsystem must register itself, usually at module_init time. This tells configfs to make the subsystem appear in the file tree.

```
struct configfs_subsystem {
```

```
diff -ru a/Documentation/filesystems/hpfs.txt b/Documentation/filesystems/hpfs.txt
--- a/Documentation/filesystems/hpfs.txt 2006-10-05 22:18:51.000000000 -0400
+++ b/Documentation/filesystems/hpfs.txt 2006-10-05 22:55:00.000000000 -0400
@@ -274,7 +274,7 @@
```

Fixed race-condition in buffer code - it is in all filesystems in Linux;

[PATCH 19-rc1] Fix typos in /Documentation : 'U-Z'

when reading device (cat /dev/hda) while creating files on it, files could be damaged

-2.02 Woraround for bug in breada in Linux. breada could cause accesses beyond
+2.02 Workaround for bug in breada in Linux. breada could cause accesses beyond end of partition

2.03 Char, block devices and pipes are correctly created

Fixed non-crashing race in unlink (Alexander Viro)

```
diff -ru a/Documentation/filesystems/ocfs2.txt b/Documentation/filesystems/ocfs2.txt
--- a/Documentation/filesystems/ocfs2.txt 2006-10-05 22:18:51.000000000 -0400
+++ b/Documentation/filesystems/ocfs2.txt 2006-10-05 23:07:22.000000000 -0400
@@ -30,7 +30,7 @@
```

Features which OCFS2 does not support yet:

- sparse files
- extended attributes
- - shared writeable mmap
- + - shared writable mmap
- loopback is supported, but data written will not be cluster coherent.

- quotas

```
diff -ru a/Documentation/filesystems/proc.txt b/Documentation/filesystems/proc.txt
--- a/Documentation/filesystems/proc.txt 2006-10-05 22:18:51.000000000 -0400
+++ b/Documentation/filesystems/proc.txt 2006-10-05 22:22:59.000000000 -0400
@@ -1220,9 +1220,9 @@
```

you probably should increase the lower_zone_protection setting.

The units of this tunable are fairly vague. It is approximately equal

-to "megabytes". So setting lower_zone_protection=100 will protect around 100
+to "megabytes," so setting lower_zone_protection=100 will protect around 100 megabytes of the lowmem zone from user allocations. It will also make
-those 100 megabytes unavailable for use by applications and by
+those 100 megabytes unavailable for use by applications and by pagecache, so there is a cost.

The effects of this tunable may be observed by monitoring

```
diff -ru a/Documentation/filesystems/spufs.txt b/Documentation/filesystems/spufs.txt
--- a/Documentation/filesystems/spufs.txt 2006-10-05 22:18:51.000000000 -0400
+++ b/Documentation/filesystems/spufs.txt 2006-10-05 22:53:50.000000000 -0400
@@ -210,7 +210,7 @@
```

/signal2

The two signal notification channels of an SPU. These are read-write files that operate on a 32 bit word. Writing to one of these files

- triggers an interrupt on the SPU. The value writting to the signal
 - + triggers an interrupt on the SPU. The value writing to the signal
- files can be read from the SPU through a channel read or from host user space through the file. After the value has been read by the SPU, it is reset to zero. The possible operations on an open signal1 or sig-

```
diff -ru a/Documentation/hrtimers.txt b/Documentation/hrtimers.txt
--- a/Documentation/hrtimers.txt 2006-10-05 22:18:51.000000000 -0400
+++ b/Documentation/hrtimers.txt 2006-10-05 22:26:52.000000000 -0400
@@ -30,7 +30,7 @@
```

necessitate a more complex handling of high resolution timers, which

[PATCH 19-rc1] Fix typos in /Documentation : 'U-Z'

in turn decreases robustness. Such a design still led to rather large timing inaccuracies. Cascading is a fundamental property of the timer – wheel concept, it cannot be 'designed out' without inevitably + wheel concept, it cannot be 'designed out' without inevitably degrading other portions of the timers.c code in an unacceptable way.

– the implementation of the current posix-timer subsystem on top of
diff –ru a/Documentation/ide.txt b/Documentation/ide.txt

```
--- a/Documentation/ide.txt 2006-10-05 22:18:52.000000000 -0400
+++ b/Documentation/ide.txt 2006-10-05 22:33:55.000000000 -0400
@@ -390,5 +390,5 @@
```

Wed Apr 17 22:52:44 CEST 2002 edited by Marcin Dalecki, the current maintainer.

–Wed Aug 20 22:31:29 CEST 2003 updated ide boot options to current ide.c

+Wed Aug 20 22:31:29 CEST 2003 updated ide boot options to current ide.c comments at 2.6.0-test4 time. Maciej Soltysiak <solt@xxxxxxxxxxxxxxxxxxx>

diff –ru a/Documentation/input/atarikbd.txt b/Documentation/input/atarikbd.txt

```
--- a/Documentation/input/atarikbd.txt 2006-10-05 22:18:51.000000000 -0400
+++ b/Documentation/input/atarikbd.txt 2006-10-05 22:46:46.000000000 -0400
@@ -103,7 +103,7 @@
```

5.1 Joystick Event Reporting

–In this mode, the ikbd generates a record whenever the joystick position is

+In this mode, the ikbd generates a record whenever the joystick position is changed (i.e. for each opening or closing of a joystick switch or trigger).

The joystick event record is two bytes of the form:

diff –ru a/Documentation/input/iforce-protocol.txt b/Documentation/input/iforce-protocol.txt

```
--- a/Documentation/input/iforce-protocol.txt 2006-10-05 22:18:51.000000000 -0400
+++ b/Documentation/input/iforce-protocol.txt 2006-10-05 22:32:14.000000000 -0400
@@ -23,9 +23,9 @@
```

When using USB:

OP DATA

The 2B, LEN and CS fields have disappeared, probably because USB handles frames and

–data corruption is handled or insignificant.

+data corruption is handled or insignificant.

–First, I describe effects that are sent by the device to the computer

+First, I describe effects that are sent by the device to the computer.

** Device input state

This packet is used to indicate the state of each button and the value of each

diff –ru a/Documentation/ioctl/cdrom.txt b/Documentation/ioctl/cdrom.txt

```
--- a/Documentation/ioctl/cdrom.txt 2006-10-05 22:18:51.000000000 -0400
+++ b/Documentation/ioctl/cdrom.txt 2006-10-05 22:25:29.000000000 -0400
@@ -735,7 +735,7 @@
```

Ok, this is where problems start. The current interface for

the CDRM_DISC_STATUS ioctl is flawed. It makes the false

assumption that CDs are all CDS_DATA_1 or all CDS_AUDIO, etc.

[PATCH 19-rc1] Fix typos in /Documentation : 'U-Z'

- Unfortunately, while this is often the case, it is also
+ Unfortunately, while this is often the case, it is also
very common for CDs to have some tracks with data, and some
tracks with audio. Just because I feel like it, I declare
the following to be the best way to cope. If the CD has

```
diff -ru a/Documentation/laptop-mode.txt b/Documentation/laptop-mode.txt
--- a/Documentation/laptop-mode.txt 2006-10-05 22:18:51.000000000 -0400
+++ b/Documentation/laptop-mode.txt 2006-10-05 22:51:06.000000000 -0400
@@ -699,7 +699,7 @@
```

Dax Kelson submitted this so that the ACPI acpid daemon will
kick off the laptop_mode script and run hdparm. The part that
automatically disables laptop mode when the battery is low was
-written by Jan Topinski.
+written by Jan Topinski.

-----/etc/acpi/events/ac_adapter BEGIN-----

```
event=ac_adapter
diff -ru a/Documentation/MSI-HOWTO.txt b/Documentation/MSI-HOWTO.txt
--- a/Documentation/MSI-HOWTO.txt 2006-10-05 22:18:52.000000000 -0400
+++ b/Documentation/MSI-HOWTO.txt 2006-10-05 22:31:22.000000000 -0400
@@ -219,7 +219,7 @@
```

Note that the pre-assigned IOAPIC dev->irq is valid only if the device
operates in PIN-IRQ assertion mode. In MSI-X mode, any attempt at
using dev->irq by the device driver to request for interrupt service
-may result unpredictable behavior.
+may result unpredictable behavior.

For each MSI-X vector granted, a device driver is responsible for calling
other functions like request_irq(), enable_irq(), etc. to enable

```
diff -ru a/Documentation/networking/NAPI_HOWTO.txt b/Documentation/networking/NAPI_HOWTO.txt
--- a/Documentation/networking/NAPI_HOWTO.txt 2006-10-05 22:18:52.000000000 -0400
+++ b/Documentation/networking/NAPI_HOWTO.txt 2006-10-05 22:52:48.000000000 -0400
@@ -601,7 +601,7 @@
```

5) dev->close() and dev->suspend() issues

-The driver writer neednt worry about this. The top net layer takes
+The driver writer needn't worry about this; the top net layer takes
care of it.

6) Adding new Stats to /proc

```
@@ -622,9 +622,9 @@
```

packets fast enough i.e send a pause only when you run out of rx buffers.

Note FC in itself is a good solution but we have found it to not be
much of a commodity feature (both in NICs and switches) and hence falls
-under the same category as using NIC based mitigation. Also experiments
-indicate that its much harder to resolve the resource allocation
-issue (aka lazy receiving that NAPI offers) and hence quantify its usefulness
+under the same category as using NIC based mitigation. Also, experiments
+indicate that it's much harder to resolve the resource allocation
+issue (aka lazy receiving that NAPI offers) and hence quantify its usefulness

[PATCH 19-rc1] Fix typos in /Documentation : 'U-Z'

proved harder. In any case, FC works even better with NAPI but is not necessary.

```
diff -ru a/Documentation/networking/sk98lin.txt b/Documentation/networking/sk98lin.txt
--- a/Documentation/networking/sk98lin.txt 2006-10-05 22:18:52.000000000 -0400
+++ b/Documentation/networking/sk98lin.txt 2006-10-05 22:40:46.000000000 -0400
@@ -346,7 +346,7 @@
```

depending on the load of the system. If the driver detects that the system load is too high, the driver tries to shield the system against too much network load by enabling interrupt moderation. If – at a later – time – the CPU utilization decreases again (or if the network load is + time – the CPU utilization decreases again (or if the network load is negligible) the interrupt moderation will automatically be disabled.

Interrupt moderation should be used when the driver has to handle one or more

```
diff -ru a/Documentation/networking/slicecom.txt b/Documentation/networking/slicecom.txt
--- a/Documentation/networking/slicecom.txt 2006-10-05 22:18:52.000000000 -0400
+++ b/Documentation/networking/slicecom.txt 2006-10-05 22:30:36.000000000 -0400
@@ -126,7 +126,7 @@
```

Though the options below are to be set on a single interface, they apply to the whole board. The restriction, to use them on 'UP' interfaces, is because the –command sequence below could lead to unpredictable results. +command sequence below could lead to unpredictable results.

```
# echo 0 >boardnum
```

```
# echo internal >clock_source
```

```
diff -ru a/Documentation/networking/wan-router.txt b/Documentation/networking/wan-router.txt
--- a/Documentation/networking/wan-router.txt 2006-10-05 22:18:52.000000000 -0400
+++ b/Documentation/networking/wan-router.txt 2006-10-05 22:41:17.000000000 -0400
@@ -412,7 +412,7 @@
```

beta3-2.1.4 Jul 2000 o X25 M_BIT Problem fix.

o Added the Multi-Port PPP

– Updated utilites for the Multi-Port PPP.

+ Updated utilities for the Multi-Port PPP.

2.1.4 Aut 2000

o In X25API:

```
@@ -450,7 +450,7 @@
```

o Keyboard Led Monitor/Debugger

– A new utility /usr/sbin/wpkbdmon uses keyboard leds

+ A new utility /usr/sbin/wpkbdmon uses keyboard leds

to convey operational statistic information of the

Sangoma WANPIPE cards.

NUM_LOCK = Line State (On=connected, Off=disconnected)

```
diff -ru a/Documentation/pnp.txt b/Documentation/pnp.txt
--- a/Documentation/pnp.txt 2006-10-05 22:18:52.000000000 -0400
+++ b/Documentation/pnp.txt 2006-10-05 22:29:15.000000000 -0400
```

[PATCH 19-rc1] Fix typos in /Documentation : 'U-Z'

@@ -184,7 +184,7 @@

Please note that the character 'X' can be used as a wild card in the function portion (last four characters).

ex:

```
- /* Unkown PnP modems */  
+ /* Unknown PnP modems */  
{ "PNPCXXX", UNKNOWN_DEV },
```

Supported PnP card IDs can optionally be defined.

```
diff -ru a/Documentation/robust-futex-ABI.txt b/Documentation/robust-futex-ABI.txt  
--- a/Documentation/robust-futex-ABI.txt 2006-10-05 22:18:52.000000000 -0400  
+++ b/Documentation/robust-futex-ABI.txt 2006-10-05 22:44:41.000000000 -0400  
@@ -170,7 +170,7 @@
```

- 1) the 'head' pointer or an subsequent linked list pointer is not a valid address of a user space word
- 2) the calculated location of the 'lock word' (address plus - 'offset') is not the valid address of a 32 bit user space word

3) if the list contains more than 1 million (subject to future kernel configuration changes) elements.

```
diff -ru a/Documentation/s390/Debugging390.txt b/Documentation/s390/Debugging390.txt  
--- a/Documentation/s390/Debugging390.txt 2006-10-05 22:18:50.000000000 -0400  
+++ b/Documentation/s390/Debugging390.txt 2006-10-05 22:43:17.000000000 -0400  
@@ -846,9 +846,9 @@
```

instead if the code isn't compiled -g, as it is much faster:

```
objdump --disassemble-all --syms vmlinux > vmlinux.lst
```

-As hard drive space is valuble most of us use the following approach.

+As hard drive space is valuable, most of us use the following approach.

- 1) Look at the emitted psw on the console to find the crash address in the kernel.
- 2) Look at the file System.map (in the linux directory) produced when building the kernel to find the closest address less than the current PSW to find the offending function.
- +2) Look at the file System.map (in the linux directory) produced when building the kernel to find the closest address less than the current PSW to find the offending function.

3) use grep or similar to search the source tree looking for the source file

```
diff -ru a/Documentation/scsi/ibmmca.txt b/Documentation/scsi/ibmmca.txt  
--- a/Documentation/scsi/ibmmca.txt 2006-10-05 22:18:50.000000000 -0400  
+++ b/Documentation/scsi/ibmmca.txt 2006-10-05 22:48:09.000000000 -0400  
@@ -461,7 +461,7 @@
```

This needs the RD-Bit to be disabled on IM_OTHER_SCSI_CMD_CMD which allows data to be written from the system to the device. It is a

necessary step to be allowed to set blocksize of SCSI-tape-drives and

- the tape-speed, without confusing the SCSI-Subsystem.

+ the tape-speed, without confusing the SCSI-Subsystem.

2) The recognition of a tape is included in the check_devices routine.

This is done by checking for TYPE_TAPE, that is already defined in

the kernel-scsi-environment. The markup of a tape is done in the

```
diff -ru a/Documentation/scsi/ncr53c8xx.txt b/Documentation/scsi/ncr53c8xx.txt  
--- a/Documentation/scsi/ncr53c8xx.txt 2006-10-05 22:18:50.000000000 -0400  
+++ b/Documentation/scsi/ncr53c8xx.txt 2006-10-05 22:36:37.000000000 -0400
```

[PATCH 19-rc1] Fix typos in /Documentation : 'U-Z'

@@ -115,7 +115,7 @@

<ftp://ftp.symbios.com/>

-Usefull SCSI tools written by Eric Youngdale are available at tsx-11:
+Useful SCSI tools written by Eric Youngdale are available at tsx-11:

<ftp://tsx-11.mit.edu/pub/linux/ALPHA/scsi/scsiinfo-X.Y.tar.gz>

<ftp://tsx-11.mit.edu/pub/linux/ALPHA/scsi/scsidev-X.Y.tar.gz>

diff -ru a/Documentation/sound/alsa/Audigy-mixer.txt b/Documentation/sound/alsa/Audigy-mixer.txt
--- a/Documentation/sound/alsa/Audigy-mixer.txt 2006-10-05 22:18:51.000000000 -0400
+++ b/Documentation/sound/alsa/Audigy-mixer.txt 2006-10-05 22:57:37.000000000 -0400

@@ -6,7 +6,7 @@

The EMU10K2 chips have a DSP part which can be programmed to support various ways of sample processing, which is described here.

-(This acticle does not deal with the overall functionality of the

+(This article does not deal with the overall functionality of the EMU10K2 chips. See the manuals section for further details.)

The ALSA driver programs this portion of chip by default code

diff -ru a/Documentation/sound/alsa/SB-Live-mixer.txt b/Documentation/sound/alsa/SB-Live-mixer.txt
--- a/Documentation/sound/alsa/SB-Live-mixer.txt 2006-10-05 22:18:51.000000000 -0400
+++ b/Documentation/sound/alsa/SB-Live-mixer.txt 2006-10-05 22:57:50.000000000 -0400

@@ -5,7 +5,7 @@

The EMU10K1 chips have a DSP part which can be programmed to support various ways of sample processing, which is described here.

-(This acticle does not deal with the overall functionality of the

+(This article does not deal with the overall functionality of the EMU10K1 chips. See the manuals section for further details.)

The ALSA driver programs this portion of chip by default code

diff -ru a/Documentation/uml/UserModeLinux-HOWTO.txt
b/Documentation/uml/UserModeLinux-HOWTO.txt
--- a/Documentation/uml/UserModeLinux-HOWTO.txt 2006-10-05 22:18:48.000000000 -0400
+++ b/Documentation/uml/UserModeLinux-HOWTO.txt 2006-10-05 23:07:57.000000000 -0400

@@ -1477,7 +1477,7 @@

- Making it world-writable looks bad, but it seems not to be
+ Making it world-writable looks bad, but it seems not to be exploitable as a security hole. However, it does allow anyone to create useless tap devices (useless because they can't configure them), which is a DOS attack. A somewhat more secure alternative would to be

-

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxxx

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

[PATCH 19-rc1] Fix typos in /Documentation : 'U-Z'

Please read the FAQ at <http://www.tux.org/lkml/>