

LINUX for S/390



LINUX[®] for S/390[®]

Device Drivers and Installation Commands

LINUX for S/390



LINUX[®] for S/390[®]

Device Drivers and Installation Commands

Third Edition – (15 December 2000)

This edition applies to the third release of the LINUX for S/390 kernel 2.2.16 patch (made in May 2000) and to all subsequent releases and modifications until otherwise indicated in new editions.

© Copyright International Business Machines Corporation 2000. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Summary of changes v

Edition 3 changes. v

Edition 2 changes. v

About this book. vii

How this book is organized. vii

Who should read this book. vii

Assumptions. vii

Part 1. Common device support . . . 1

Part 2. LINUX for S/390 Open-Source device drivers 3

Chapter 1. LINUX for S/390 DASD device driver 5

DASD overview 5

Partitioned DASD 6

DASD features. 7

DASD kernel parameter syntax 7

DASD kernel example 9

DASD module parameter syntax. 9

DASD module example 10

DASD – Preparing for use 10

1) Low level format. 10

2) Make a file system 10

DASD restrictions 11

Chapter 2. LINUX for S/390 VM minidisk device driver 13

VM minidisk features 13

VM minidisk kernel parameter syntax 13

VM minidisk kernel example 13

VM minidisk – Preparing disks. 14

Chapter 3. LINUX for S/390 XPRAM device driver 17

XPRAM features. 17

XPRAM kernel parameter syntax 17

XPRAM kernel example 18

XPRAM module parameter syntax. 18

XPRAM module example. 18

Chapter 4. LINUX for S/390 CTC/ESCON device driver 19

CTC/ESCON features 19

CTC/ESCON kernel parameter syntax 19

CTC/ESCON kernel example 21

CTC/ESCON module parameter syntax 21

CTC/ESCON module example 22

CTC/ESCON – Preparing the connection 22

CTC/ESCON – Recovery procedure after a crash. 24

Chapter 5. LINUX for S/390 IUCV device driver 25

IUCV features 25

IUCV kernel parameter syntax 26

IUCV kernel parameter example 26

IUCV module parameter syntax 26

IUCV module parameter example 27

IUCV – Preparing the connection 27

IUCV – Further information 29

IUCV restrictions 29

Chapter 6. LINUX for S/390 Console device drivers. 31

Console features. 31

Console kernel parameter syntax 31

Console kernel examples 32

Using the console 32

Console special characters 32

Console 3270 emulation 33

Console – Use of VInput 33

Console limitations 34

Part 3. LINUX for S/390 Object-Code-Only device drivers . . 35

Chapter 7. LINUX for S/390 LCS Device Driver 37

LCS features 37

LCS kernel parameter syntax 38

LCS kernel parameter example 39

LCS module parameter syntax 39

LCS module parameter example 40

LCS restrictions 40

LCS limitations 40

LCS – Common set up problem 40

Chapter 8. LINUX for S/390 Gigabit Ethernet device driver 43

GbE features 43

GbE module parameter syntax 43

GbE examples 45

1: Basic configuration 45

2: Router configuration 46

GbE – Preparing the connection 46

GbE device recognition 47

GbE restrictions 48

GbE queuing 48

GbE background – QDIO. 50

Part 4. Installation commands and parameters 51

Chapter 9. Useful LINUX commands . . . 53

dasdfmt - Format a DASD	54
ifconfig - Configure a network interface	56
insmod - Load a module into the LINUX kernel	60
mke2fs - Create a file system	62
silos - Make DASD bootable	63

Chapter 10. Kernel parameters 65

ipldelay	66
maxcpus	67
mem.	68
noinitrd	69
ramdisk_size	70
ro	71
root	72
vmhalt	73

Chapter 11. Overview of the parameter line file. 75

Parameters	76
----------------------	----

Part 5. Appendixes. 79

Appendix A. Reference information . . 81

LCS module parameter syntax	81
LCS kernel parameter syntax	81
Gigabit Ethernet driver command syntax	81
LINUX for S/390 Device numbers	82

Appendix B. Kernel building 83

Building the kernel	83
Preliminary steps	84
Configuring the parameters	85
Checking the configuration	86
Checking the dependencies	86
Compiling the kernel	86
Installing the modules	87
Finishing off	87
Using 'config' or 'oldconfig'	88
Sample output listing	88
Cross-reference to configuration options	90
Using 'menuconfig'	91
File handling	91
Main menu	93
Code maturity level option	94
Processor type and features	94
Loadable module support	95
General setup	95
S/390 block device drivers	96

S/390 network device support	97
S/390 terminal and console options	97
Networking options	98
QoS and/or Fair queueing	99
Filesystems	99
Network file systems	100
Partition types	100
Kernel hacking	101
Load an alternate configuration file	101
Save configuration to an alternate file	102
Exit 'menuconfig'	102
Kernel parameter options	103
IEEE FPU emulation	103
Built-in IPL record support	104
IPL method generated into head.S	104
Support for VM minidisk	104
Support for VM minidisk synchronous read-write	104
Support for DASD devices	105
Support for ECKD disks	105
Support for FBA disks	105
Support for DIAG access to CMS reserved minidisk	106
XPRAM device support	106
CTC/ESCON device support	106
IUCV device support	107
Support for 3215 line mode terminal	107
Support for console output on 3215 line mode terminal	107
Support for hardware console	108
Console output on hardware console	108

Glossary 109

Notices 113

Trademarks	114
----------------------	-----

International License Agreement for Non-Warranted Programs 115

GNU General Public Licence, Version 2, June 1991 117

Preamble	117
GNU General Public Licence: Terms and conditions for copying, distribution and modification	118
How to Apply These Terms to Your New Programs	121

Index 123

Summary of changes

This revision contains changes to support the LINUX for S/390 kernel loadable module for the LINUX kernel version 2.2.16.

Edition 3 changes

New Information

- Gigabit Ethernet driver restriction

This revision also includes maintenance and editorial changes. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change.

Edition 2 changes

New Information

- CTC/ESCON VM channel subset selection for TCP/IP

Changed Information

- CTC/ESCON module parameter syntax
- VM Minidisk driver revisions
- 'mem' parameter additional option
- Console parameter change for P/390

About this book

This book describes the drivers available to LINUX for the control of S/390 devices and attachments.

The drivers described herein have been developed with version 2.2.16 of the LINUX kernel. If you are using a later version of the kernel, the kernel parameters may be different to those described in this document.

For more specific information about the device driver structure, see the documents in the kernel source tree at `...linux/Documentation/s390`.

When you have installed Linux including the kernel sources this path will be on your machine. Typically: `/usr/src/linux/Documentation/s390`.

How this book is organized

The first part of this book contains general information relevant to all LINUX for S/390 device drivers.

Parts two and three consist of chapters specific to individual device drivers. (Part two describes the drivers supplied as open-source; part three describes the drivers supplied as object-code-only (OCO).)

Part four contains information on the LINUX and S/390 commands and parameters used in installing.

These chapters are followed by a reference section containing summaries of the command syntax of the drivers, a glossary and an index.

Who should read this book

This book is intended for :

- System administrators who wish to configure a LINUX for S/390 system

Assumptions

The following general assumptions are made about your background knowledge:

- You have an understanding of LINUX and S/390 terminology.
- You are familiar with LINUX device driver software.
- You have an understanding of basic computer architecture, operating systems, and programs.
- You are familiar with the S/390 devices attached to your system. (S/390 knowledge should not be required, as the code specific to the S/390 hardware is provided by IBM.)

Part 1. Common device support

Before LINUX for S/390 can use a device the associated device driver must be available to the LINUX kernel. This can be achieved either by compiling the device driver into the kernel or by invoking the driver as a module. The options for each driver are shown in the following table:

Device driver	Kernel	Module
DASD	yes	yes
VM minidisk	yes	no
XPRAM	yes	yes
CTC/ESCON	yes	yes
IUCV	yes	no
Hardware console	yes	no
3215 console	yes	no
LCS	no	yes
Gigabit Ethernet	no	yes

A description of how to build the kernel including device drivers is given in “Appendix B. Kernel building” on page 83.

The parameters for the kernel resident device drivers are held in the parameter line file which is created during the installation of LINUX .

- If you are using an LPAR or native installation this is parameter -p in the silo parameter file.
- For a VM installation, include the parameter in the PARM LINE A file.

For the format of this file see “Chapter 11. Overview of the parameter line file” on page 75.

Drivers which are not kernel resident are loaded into LINUX with their parameters by means of the insmod command. See “insmod - Load a module into the LINUX kernel” on page 60 for the syntax.

Because the ESA/390 architecture differs from that used by the Intel PC and other machines the I/O concepts used by S/390 device drivers are also different.

LINUX was originally designed for the Intel PC architecture which uses two cascaded 8259 programmable interrupt controllers (PIC) that allow a maximum of 15 different interrupt lines. All devices attached to that type of system share those 15 interrupt levels (or IRQs). In addition, the bus systems (ISA, MCA, EISA, PCI, etc.) might allow shared interrupts, different polling methods or DMA processing.

Unlike other hardware architectures, ESA/390 implements a channel subsystem that provides a unified view of the devices attached to the system. Although a large variety of peripheral attachments are defined for the ESA/390 architecture, they are all accessed in the same manner using I/O interrupts. Each device attached to the system is uniquely identified by a subchannel, and the ESA/390 architecture allows up to 64,000 devices to be attached.

To avoid the introduction of a new I/O concept to the common LINUX code, LINUX for S/390 preserves the IRQ concept and semantically maps the ESA/390 subchannels to LINUX as IRQs. This allows LINUX for S/390 to support up to 64K different IRQs, each representing a unique device.

The unified I/O access method incorporated in LINUX for S/390 allows the operating system to implement all of the hardware I/O attachment functionality that each device driver would otherwise have to provide itself. A common I/O device driver is provided which uses a functional layer to provide a generic access method to the hardware. The driver comprises a set of I/O support routines, some of which are common LINUX interfaces, while others are LINUX for S/390 specific:

get_dev_info()

Allows a device driver to find out what devices are attached (visible) to the system, and to determine their current status.

request_irq()

Assigns the ownership of a specific device to a device driver.

free_irq()

Releases the ownership of a specific device.

disable_irq()

Prevents a specific device from presenting interrupts to the device driver.

enable_irq()

Allows a device to present I/O interrupts to the device driver.

do_I0()

Initiates an I/O request.

halt_I0()

Terminates the I/O request that is currently being processed by the device.

do_IRQ()

This is an interrupt pre-processing routine that is called by the interrupt entry routine whenever an I/O interrupt is presented to the system. The `do_I0()` routine determines the interrupt status and calls the device specific interrupt handler according to the rules (flags) defined by `do_I0()`.

More information on these commands can be found in the Linux source directory, `.../Documentation/s390/cds.txt`

Part 2. LINUX for S/390 Open-Source device drivers

The open-source device drivers are:

- “Chapter 1. LINUX for S/390 DASD device driver” on page 5
- “Chapter 2. LINUX for S/390 VM minidisk device driver” on page 13
- “Chapter 3. LINUX for S/390 XPRAM device driver” on page 17
- “Chapter 4. LINUX for S/390 CTC/ESCON device driver” on page 19
- “Chapter 5. LINUX for S/390 IUCV device driver” on page 25
- “Chapter 6. LINUX for S/390 Console device drivers” on page 31

Chapter 1. LINUX for S/390 DASD device driver

DASD overview

The DASD device driver in LINUX for S/390 takes care of all real or emulated DASD (Direct Access Storage Device) that can be attached to an S/390 system. The class of devices named DASD includes a variety of physical media, on which data is organized in blocks and/or records which can be accessed (read or written) in random order.

Traditionally these devices are attached to a control unit connected to an S/390 I/O channel. In modern systems these have been largely replaced by emulated DASD, such as the internal disks of the Multiprise family, the volumes of the RAMAC virtual array, or the volumes of the Enterprise Storage Server. These are completely virtual representations of DASD in which the identity of the physical device is hidden.

The LINUX for S/390 DASD device driver can only handle devices which are represented as volumes attached to the S/390 system by a control unit.

The driver can either be statically built into the kernel or loaded during run time as a module.

The DASD device driver is capable of accessing an arbitrary number of devices. The default major number for DASD (94) can only address 64 DASD (see below), so additional major numbers (typically descending from 254) are allocated dynamically at initialization or run time. The range of major numbers available in the dynamic allocation pool is in practice the only limit to the number of DASD accessible.

Each DASD configured to the system uses 4 minor numbers.

- The first minor number always represents the entire device, including IPL and label records.
- The remaining three minor numbers represent partitions of the device as defined in the partition table.

The DASD device driver has a built in naming scheme for DASD according to Table 1. (You can override the built in scheme by creating customized nodes in the LINUX /dev/ subdirectory.) These names are sufficient to access 18278 devices, but a LINUX 2.2 or 2.4 system is restricted to 256 major numbers and 64 blocks of 4 minor numbers, giving a maximum of 16384 DASD even if no numbers are used for other types of device. Every major number used for other devices reduces the maximum number of DASD by 64.

Table 1. DASD naming convention

Names	Number
dasda – dasdz	26
dasdaa – dasdzz	676
dasdaaa – dasdzzz	17576
Sum:	18278

General DASD nodes have the format *dasdX*, or *dasdXP*, where *X* is a letter denoting the device, and *P* is a number denoting the partition on that device. The first form, *dasdX*, is used to address the entire disk. The second, *dasdXP*, is used to address the partitions on this device.

For example `/dev/dasda` refers to the whole of the first disk in the system and `/dev/dasda1` to the first partition on that disk.

They are typically created by:

```
mknod -m 660 /dev/dasda b 94 0
mknod -m 660 /dev/dasda1 b 94 1
mknod -m 660 /dev/dasda2 b 94 2
mknod -m 660 /dev/dasda3 b 94 3
mknod -m 660 /dev/dasdb b 94 4
mknod -m 660 /dev/dasdb1 b 94 5
....
```

Partitioned DASD

The DASD device driver is embedded into the LINUX generic support for partitioned disks. This implies that you can have any kind of partition table known to LINUX on your DASD, such as the MSDOS or Amiga partition scheme. However none of the partition schemes built in to LINUX to support platforms other than S/390 will preserve S/390 IPL and label records.

'IBM label' partition scheme:

To ensure compatibility to other S/390 operating systems the IBM-label partition scheme has been added to LINUX . This scheme currently supports LNX (LINUX) and CMS (VM/ESA) labelled disks, as well as unlabeled disks which are treated equivalently to LNX-labelled disks. The disk layout of the different types is shown in Figure 1.

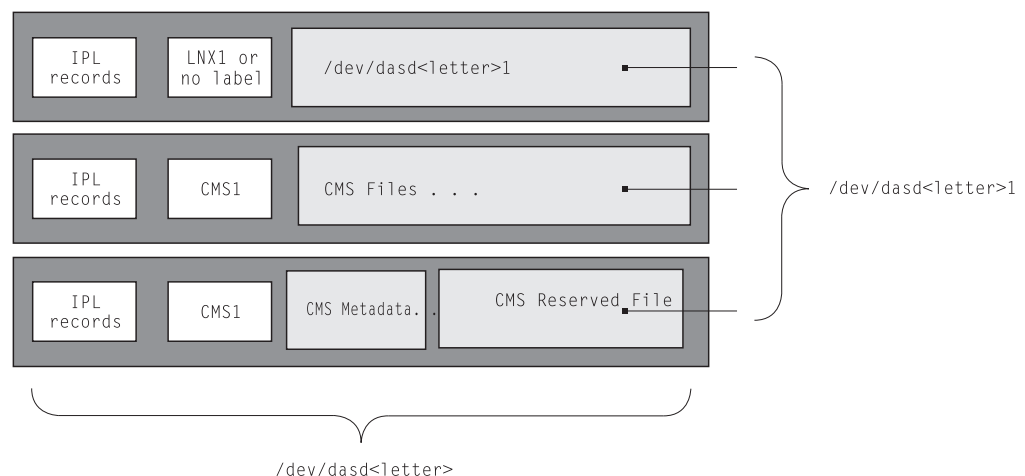


Figure 1. Partition scheme for LNX and CMS labelled disks

The first of these examples shows a disk in an LPAR or native mode, or a full pack minidisk (dedicated DASD) in VM. The second and third examples are VM specific.

LNX1 labelled disk or non-labelled volume:

These disks are implicitly reserved for use by LINUX . The disk layout reserves the IPL and label records for access through the 'entire disk' device. All remaining records are grouped into the first partition.

CMS1 labelled disk:

Handling of these disks depends on the content of the CMS filesystem. If the volume contains a CMS filesystem it will be treated equivalently to a LNX labelled volume. If the volume is a CMS reserved volume ¹ the CMS reserved file is represented by the first and only partition. IPL and label records as well as the metadata of the CMS filesystem are reserved for access through the 'entire disk' device.

DASD features

The DASD device driver can access devices according to Table 2 by its built in CCW interface.

*Table 2. Supported devices. '**' signifies any digit.*

<u>Device format</u>	<u>Control unit type/model</u>	<u>Device type/model</u>
ECKD	3990(2105)/**	3380/**
(Extended Count Key Data)	3990(2105)/**	3390/**
	9343/**	9345/**
FBA	6310/??	9336/??
(Fixed Block Access)	3880/**	3370/**

In addition under a guest operating system in VM/ESA any DASD device supported by VM/ESA is also supported by LINUX for S/390 by accessing the device using the DIAG250 command.

The DASD device driver is also known to work with these devices:

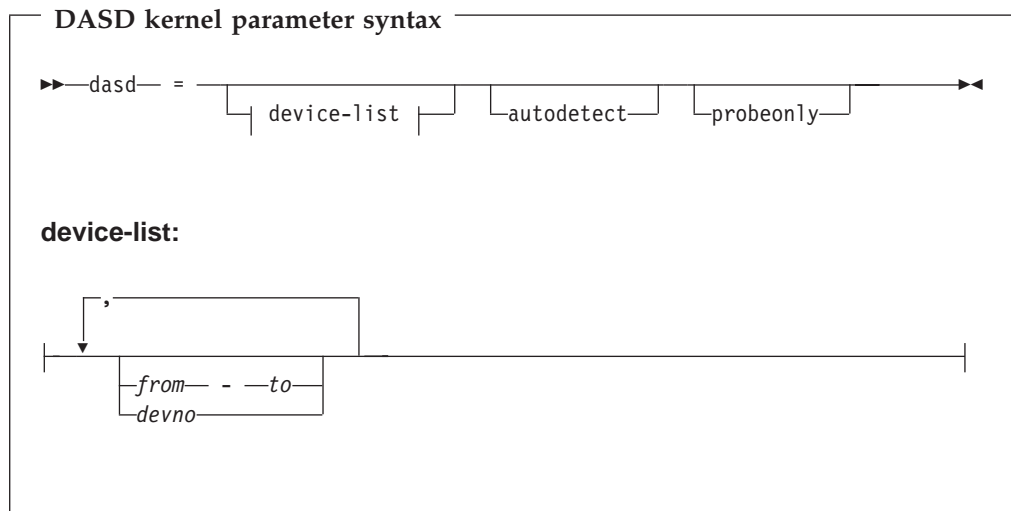
- Multiprise internal disks
- RAMAC
- RAMAC RVA
- Enterprise Storage Server (Seascape) virtual ECKD-type disks

We currently implement one partition per volume, which is the whole volume, skipping the first blocks according to the scheme outlined in Figure 1 on page 6.

DASD kernel parameter syntax

The DASD driver is configured by a kernel parameter added to the parameter line:

1. CMS reserved volume means a volume that has been reserved by a 'CMS RESERVE fn ft fm' command.



where:

autodetect

causes the driver to consider any device operational at the time of IPL as a potential dasd and allocate a device number for it. Nevertheless none of the devices which are not DASD, or not accessible by one of the DASD disciplines known to the kernel, will be accessible as a DASD. Any 'open' request on such a device will return ENODEV. In /proc/dasd/devices these devices will be flagged 'unknown'.

probeonly

causes the DASD device driver to reject any 'open' syscall with EPERM.

autodetect,probeonly

behaves in the same way as above, but additionally all devices which are accessible as DASD will refuse to be opened, returning EPERM. This setting is the default, if no 'dasd=...' parameter is given in the command line or in the module parameter.

device-list

defines a range of DASD. This is the form of the 'dasd=...' parameter to be used when you know exactly which devices you want to access with LINUX .

from-to defines the first and last DASD in a range. All DASD devices with addresses in the range are selected. It is not necessary for the *from* and *to* addresses to correspond to actual DASD.

devno defines a single DASD address.

The DASD addresses must be given in hexadecimal notation with or without a leading 0x, for example 0x191 or 5a10.

If you supply one or more kernel parameters *dasd=device-list1,dasd=device-list2,...*, the devices are processed in order of appearance in the parameter line. Devices are ignored if they are unknown to the machine, non-operational, or set off-line.²

² Currently there is no check for duplicate occurrences of the same device number.

If autodetection is turned on a DASD device is allocated in LINUX for every device operational at the time of initialization of the driver, in order of ascending subchannel numbers.

Note that the autodetection option may cause confusing results if you change your I/O configuration between two IPLs, or if you are running as a guest operating system in VM/ESA, because the devices might appear with different names (major/minor combinations) in the new IPL .

DASD kernel example

```
dasd=192-194,5a10
```

This reserves major/minor numbers and nodes as follows:

```
94 0 /dev/dasda - for the entire device 192
94 1 /dev/dasda1 - first partition on 192
94 2 /dev/dasda2 - reserved (not used)
94 3 /dev/dasda3 - reserved (not used)

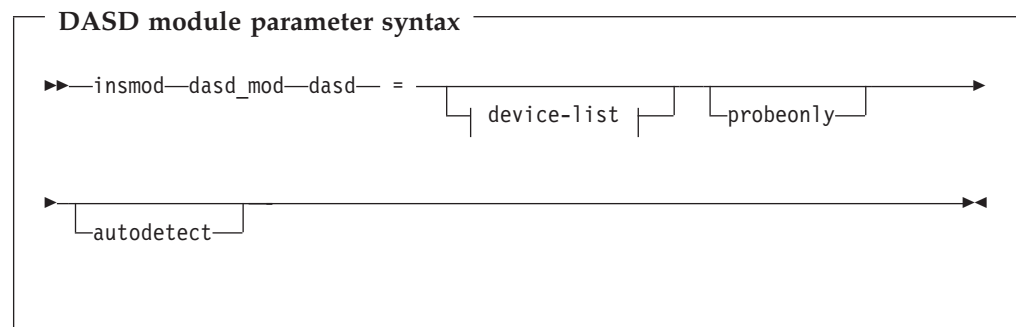
94 4 /dev/dasdb - for the entire device 193
94 5 /dev/dasdb1 - first partition on 193
94 6 /dev/dasdb2 - reserved (not used)
94 7 /dev/dasdb3 - reserved (not used)

94 8 /dev/dasdc - for the entire device 194
94 9 /dev/dasdc1 - first partition on 194
94 10 /dev/dasdc2 - reserved (not used)
94 11 /dev/dasdc3 - reserved (not used)

94 12 /dev/dasdd - for the entire device 5a10
94 13 /dev/dasdd1 - first partition on 5a10
94 14 /dev/dasdd2 - reserved (not used)
94 15 /dev/dasdd3 - reserved (not used)
```

DASD module parameter syntax

The following are the DASD driver module parameters:



where:

dasd_mod

is the name of the device driver module

dasd

is the start of the parameters

and all other parameters are the same as the dasd kernel parameters.

DASD module example

```
insmod dasd_mod dasd=192-194,5a10
```

The details are the same as “DASD kernel example” on page 9.

DASD – Preparing for use

1) Low level format

Before using an ECKD type DASD as a LINUX for S/390 disk the device must be formatted. This should be done from LINUX for S/390 by issuing an `ioctl` called `BIODASDFORMAT` on the file descriptor of the opened volume `/dev/dasd<letter>`. The utility `dasdfmt` is provided as an interface to this `ioctl` with additional checking.

Caution: Using `dasdfmt` or the raw `ioctl` can potentially destroy your running LINUX for S/390 system, forcing you to reinstall from scratch.

See the help given by `dasdfmt -help` and “`dasdfmt` - Format a DASD” on page 54 for further information. The `dasdfmt` utility calls several processes sequentially. Take care to allow sufficient time for each process to end before attempting to enter an additional command.

We recommend you set `blksize` to 1024 or higher (ideally 4096) because the `ext2fs` file system uses 1KB blocks and 50% of capacity will be unusable if the DASD `blocksize` is 512 bytes.

The formatting process can take a long time (hours) for large DASD.

2) Make a file system

Before using a DASD as a LINUX for S/390 data disk, you must create a file system on it. (A DASD for use as a swap device or paging space only needs to be defined as such.) Using `mk__fs` (replacing `__` with the appropriate identifier for the file system – for example use `mke2fs` for an `ext2` file system) you can create the file system of your choice on that volume or partition .

It is recommended that you build your file system on the first partitions of the DASD (`/dev/dasda1`, `/dev/dasdb1`, and so on), rather than the whole volume. This has a cost of 3 blocks of disk space, but it will allow you to introduce a real partition table on the device without losing access your data.

Note that the `blocksize` of the file system must be larger than or equal to the `blocksize` given to the `dasdfmt` command. It is recommended that the two `blocksize` values are equal.

You must have `CONFIG_DASD`, `CONFIG_DASD_ECKD`, `CONFIG_DASD_FBA` and `CONFIG_DASD_MDSK` enabled in the configuration of your current kernel to access IBM DASDs.

DASD restrictions

- Note that the `dasdfmt` utility can only format volumes containing a standard record zero on all tracks. If your disk does not fulfill this requirement (for example if you re-use an old volume, or access a brand new disk or one having an unknown history), you should additionally use a device support facility such as ICKDSF (in OS/390, VM, VSE or stand-alone) before doing the `dasdfmt` for the low-level format.
- Because the full set of S/390 error recovery actions are not implemented yet, you may encounter problems if you use DASDs with faulty tracks or records, particularly old 3380 and 3390 devices.
- The size of any swap device or file may not exceed 2 GB. Similarly, the limit for the main memory that can be defined is slightly less than 2 GB. Refer to “Chapter 3. LINUX for S/390 XPRAM device driver” on page 17 for information about the expanded memory capability.

Chapter 2. LINUX for S/390 VM minidisk device driver

Under VM it is possible to divide DASDs into logical partitions, known as minidisks. These minidisks have a unique 16 bit identification – the virtual device number. Each virtual device can be formatted from VM and used with the CMS file system. Also, it is possible to create a reserved minidisk, which appears to CMS as a single large file the size of the whole virtual device. This reserved file can be written to under CMS and accessed with a special opcode DIAG 250 as a block device.


Note: It is also possible under VM to attach an entire DASD to a LINUX guest (as a 'full pack minidisk'). A DASD attached in this manner is controlled by the LINUX for S/390 DASD device driver (see "DASD overview" on page 5) and not by the VM minidisk device driver.

VM minidisk features

A reserved minidisk must be formatted with a blocksize of either 512, 1024, 2048 or 4096 bytes. It may have any partition size up to the physical disk limit.

VM minidisk kernel parameter syntax

The VM Minidisk driver is configured by a kernel parameter added to the LINUX parameter line file (PARM LINE A):



where:

vdev virtual device number as hex number (without the leading 0x)

It is possible to have more than one `mdisk=` statement in the PARM LINE A file; the assignment to the device minor numbers will follow the order of the statements.

VM minidisk kernel example

The command:

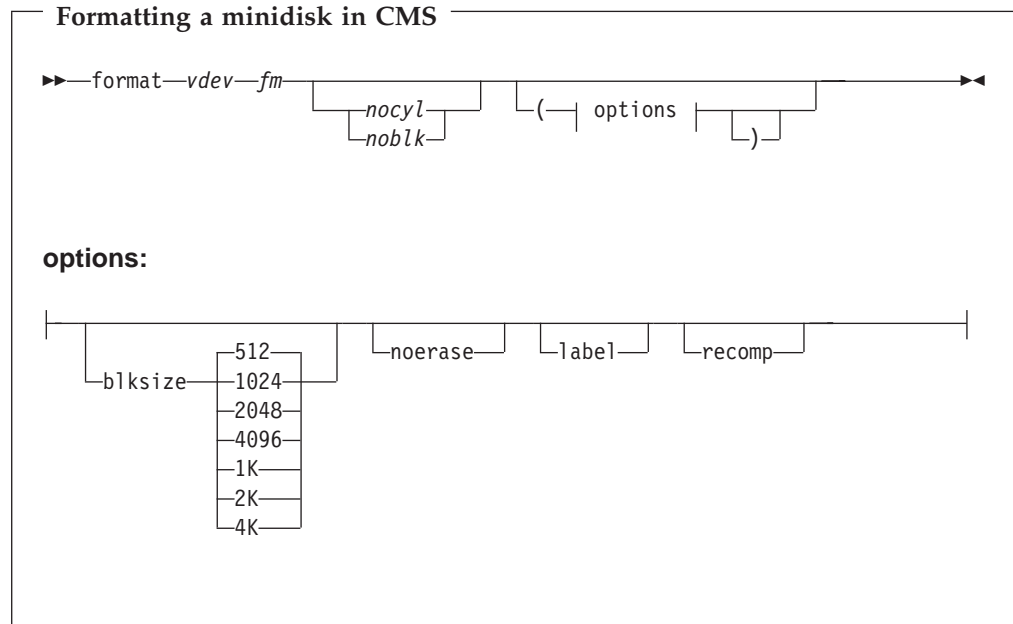
```
mdisk=193,194
```

allocates two minidisks to LINUX with device numbers 0x193 and 0x194

VM minidisk – Preparing disks

A reserved minidisk is created with the following steps in CMS:

1. Format minidisk.
2. Reserve minidisk.



Where:

vdev is the unit address (hexadecimal with no leading 0x)

fm is the CMS disk access letter (one character, A-R or T-Z)

nocyl is the number of cylinders to be allocated (non-FB-512 devices)


noblk is the number of blocks to be allocated (FB-512 devices only)

blksize
is the size of blocks to be formatted

noerase
means that the blocks are not to be cleared to zeroes (FB-512 devices only)

label is the label to be assigned to the minidisk (1 to 6 alphanumeric characters).
If this is the only parameter given then the disk is re-labelled without re-formatting.

recomp
changes the number of cylinders/blocks available on a previously formatted minidisk



Where:

filename filetype

is a valid CMS file name (each part is 1 to 8 characters)

filemode

is the CMS disk access letter specified in the format command

Example:

```
format 192 b (blksize 4096
reserve mnda mnda b
```

Making a file system

Before using a VM minidisk as a LINUX for S/390 data disk, you must create a file system on it. (A VM minidisk for use as a swap device or paging space only needs to be defined as such.) Using `mk__fs` (replacing `__` with the appropriate identifier for the file system – for example use `mke2fs` for an ext2 file system) you can create the file system of your choice on that volume or partition .

Note that the blocksize of the file system must be larger than or equal to the blocksize given to the format command. It is recommended that the two blocksize values are equal.

Chapter 3. LINUX for S/390 XPRAM device driver

The S/390 hardware supports more physical memory than can be accessed as main memory at one time. The LINUX for S/390 main memory is limited to 2 GB. However, additional memory can be declared as expanded storage. The S/390 architecture allows applications to access up to 16 TB of expanded storage and the zSeries will allow 18 EB (exabytes) (although the current hardware, S/390 and zSeries, can be equipped with at most 64 GB of real memory). The memory in the expanded storage range can be swapped in or out of the main memory in 4 KB blocks.

An IPL (boot) of LINUX for S/390 does not reset expanded storage, so it is persistent through IPLs and could be used, for example, to store diagnostic information. It is of course reset by an IML (power off/on).


The XPRAM device driver is a block device driver that enables LINUX for S/390 to access the expanded storage. Thus XPRAM can be used as a basis for fast swap devices and/or fast file systems.

XPRAM features

- Automatic detection of expanded storage.
(If expanded storage is not available, XPRAM fails gracefully with a log message reporting the lack of expanded storage.)
- Storage can be subdivided into up to 32 partitions.
- Device driver major number: 35.
- Partition minor numbers: 0 through 31.
- Hard sector size: 4096 bytes.

XPRAM kernel parameter syntax

The kernel parameter is optional. If omitted the default is to define the whole expanded storage as one partition. The syntax is:



where *number_of_partitions* defines how many partitions the expanded storage is split into. The *i*-th *partition_size* defines the size of the *i*-th partition. Blank entries are inserted if necessary to fill *number_of_partitions* values. Each size may be blank, specified as a decimal value, or a hexadecimal value preceded by 0x, and may be qualified by a magnitude:

- k or K for kilo (1024) is the default
- m or M for Mega (1024*1024)

- g or G for Giga (1024*1024*1024)

The size value multiplied by the magnitude defines the partition size in bytes. The default size is 0.

Any partition defined with a non-zero size is allocated the amount of memory specified by its size parameter.

Any remaining memory is divided as equally as possible among any partitions with a zero or blank size parameter, subject to the two constraints that blocks must be allocated in multiples of 4K and addressing constraints may leave unallocated areas of memory between partitions.

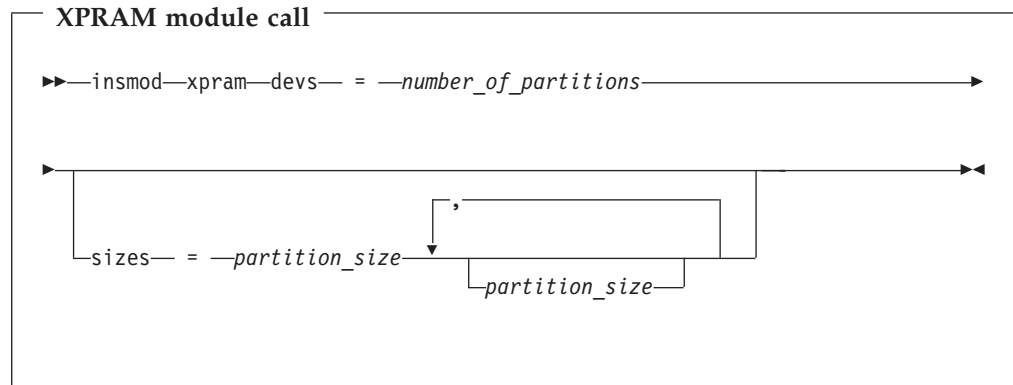
XPRAM kernel example

```
xpram_parts=4,0x800M,0,0,0x1000M
```

This allocates the extended storage into four partitions. Partition 1 has 2 GB (hex 800M), partition 4 has x 4 GB, and partitions 2 and 3 use equal parts of the remaining storage. If the total amount of extended storage was 16 GB, then partitions 3 and 4 would each have approximately 5 GB.

XPRAM module parameter syntax

If it is not included in the kernel XPRAM may be loaded as a module. The syntax of the module parameters passed to insmod differs from the kernel parameter syntax:



where:

- *partition_size* is a non-negative integer that defines the size of the partition in KB. Only decimal values are allowed and no magnitudes are accepted.

XPRAM module example

```
insmod xpram devs=4 sizes=2097152,8388608,4194304,2097152
```

This allocates a total of 16 GB of extended storage into four partitions, of (respectively) size 2 GB, 8 GB, 4 GB, and 2 GB.

Chapter 4. LINUX for S/390 CTC/ESCON device driver

A CTC connection or an ESCON connection is the typical high speed connection between mainframes. The data packages and the protocol of both connections are the same. The difference between them is the physical channel used to transfer the data.

- A CTC channel uses a parallel copper cable to connect one mainframe to another. This is similar to (though much faster than) a point-to-point connection between two PCs using a parallel cable to connect their parallel ports.
- An ESCON channel uses a fibre optic connection, and the signals are passed in sequence over the connection. This is similar to (though much faster than) a point-to-point connection between two PCs using a serial cable to connect their serial ports.

Both types of connection may be used to connect a mainframe, an LPAR, or a VM guest to another mainframe, LPAR or VM guest, where the peer LPAR or VM guest may reside on the same or on a different system.

A third type of connection is virtual CTC which is a software connection between two VM guests on the same VM system and which is faster than a physical connection.

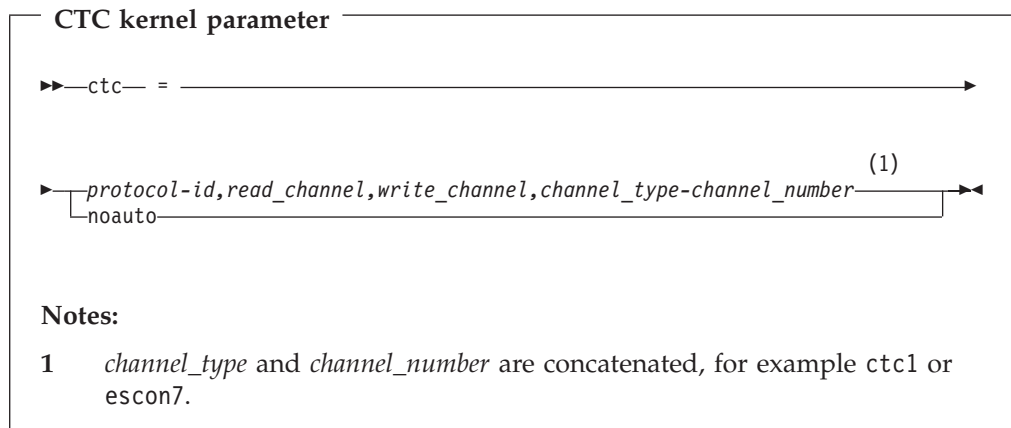
The LINUX for S/390 CTC device driver supports all three types of connection and can be used to establish a point-to-point TCP/IP connection between two LINUX for S/390 systems or between a LINUX for S/390 system and another operating system such as VM, VSE or z/OS.

CTC/ESCON features

- 8 CTC and/or ESCON connections available (in current implementation).
- Autosense mode available (the driver will pick the 16 channels with the lowest subchannel IDs).
- 256 channels to choose from. (When a CTC definition is found in the kernel parameter file the channels must be in the first 256 CTC/ESCON subchannel ID's. This value of 256 (MAX_CHANNEL_DEVICES) can be changed in the `ctc.c` file.)

CTC/ESCON kernel parameter syntax

The default for this driver is to select channels in order (automatic channel selection). If you need to use the channels in a different order, or do not want to use automatic channel selection, you can specify alternatives using the `ctc=` kernel parameter.



Note: The entire parameter is repeated (separated by spaces) for each CTC/ESCON device.

Where:

protocol-id
 must always be 0

read_channel
 is the read channel address (in hexadecimal preceded by 0x)

write_channel
 is the write channel address (in hexadecimal preceded by 0x)

channel_type
 is *ctc* or *escon*

channel_number
 is the channel number (0 to 7)

Using the *ctc=* keyword with any parameter switches automatic channel selection off.

If the only parameter given is *noauto* the CTC driver is disabled. This might be necessary, for example, if your installation uses 3271 devices or other such devices that use the CTC device type and model, but operate with a different protocol.

CTC/ESCON kernel example

For one network device (CTC):



Figure 2. Connection of two systems via CTC

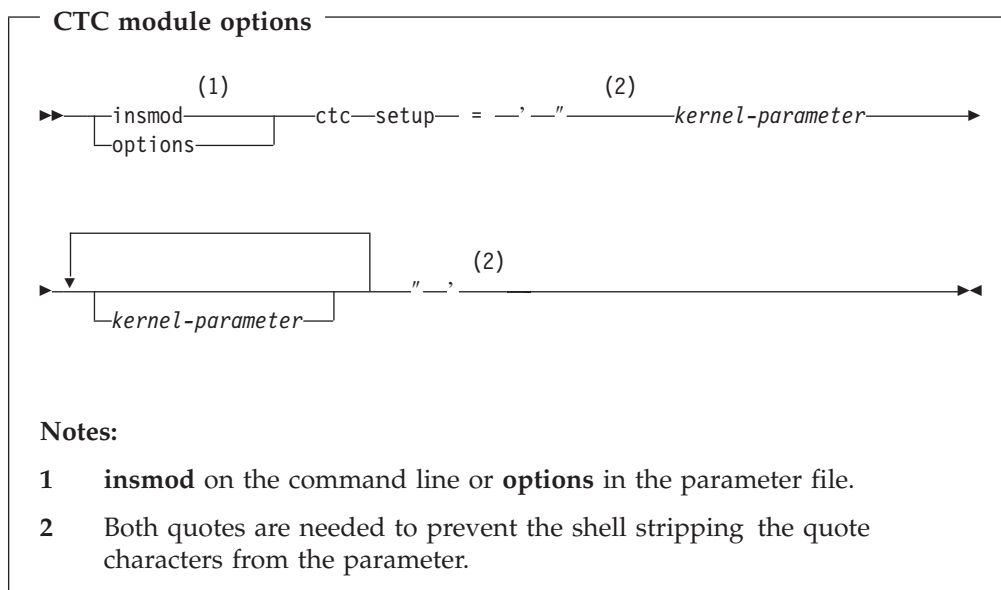
```
ctc=0,0x600,0x601,ctc0
```

Or for two network devices (CTC + ESCON):

```
ctc=0,0x601,0x600,ctc0 ctc=0,0x605,0x608,escon3
```

CTC/ESCON module parameter syntax

These parameters can be passed to the CTC/ESCON driver module by `insmod`, or can be specified in the parameter file `"/etc/modules.conf"` or `"/etc/conf.modules"` (the file name depends on the LINUX distribution).



Where:

kernel-parameter

is as defined above in "CTC/ESCON kernel parameter syntax" on page 19

Note: If the parameter line file is used the CTC driver may be loaded by typing `modprobe ctc` on the command line.

CTC/ESCON module example

For one network device (CTC):




Figure 3. Connection of two systems via CTC

Command line example:

```
insmod ctc setup='"ctc=0,0x0600,0x0601,ctc0"'
```

or

Parameter file example:

```
options ctc setup='"ctc=0,0x0600,0x0601,ctc0"'
```

Or for two network devices (CTC + ESCON):

Command line example:

```
insmod ctc setup='"ctc=0,0x0601,0x0600,ctc0 ctc=0,0x0605,0x0608,escon3"'
```

or

Parameter file example:

```
options ctc setup='"ctc=0,0x0601,0x0600,ctc0 ctc=0,0x0605,0x0608,escon3"'
```

CTC/ESCON – Preparing the connection

1. Connection

Prior to activation a channel connection is required. This can be a real or virtual connection :

- Real Channels

Connect the systems with a pair of channels to the remote system. Verify that the read channel of one is connected to the write channel of the other.

- LPAR to LPAR Channels

Select a pair of channels on each system. Verify that the read channel of one is connected to the write channel of the other and vice-versa.

- VM Channels

- a. Obtain a subnet from your TCP/IP communications staff. It is important that the subnet used by your LINUX guests not be the same as that used by VM on the LAN. It is a separate network and should be treated as such.

- b. Take one address from that subnet and assign it to VM.

- c.

Define two virtual channels to your user ID. The channels may be defined in the VM User Directory using directory control SPECIAL statements, for example:


```
special 0c04 ctca
special 0c05 ctca
```

or by using the CP commands:

```
define ctca as 0c04
define ctca as 0c05
```

from the console of the running CMS machine (preceded by #CP if necessary), or from an EXEC file (such as PROFILE EXEC A).

- d. Add the necessary VM TCP/IP routing statements (BsdRoutingParms or Gateway). Use an MTU size of 9216 and a point-to-point host route (subnet mask 255.255.255.255). If you use dynamic routing, but do not wish to run routed or gated on LINUX, update the VM ETC GATEWAYS file to include "permanent" host entries for each LINUX guest.
- e. Bring these updates on-line by using OBEYFILE or by recycling TCP/IP and/or ROUTED as needed.
- f. Add the statement
IUCV ALLOW

to your VM user directory entry.

Connect the virtual channels to the channels of the VM TCP/IP target user ID. You must couple the LINUX read channel to the VM TCP/IP write channel and vice versa. The coupling can be done with the following CP commands (following the previous example)

```
couple 0c04 to tcpip 0c05
couple 0c05 to tcpip 0c04
```

The VM TCP/IP channel numbers depend on the customisation on the remote side. In this example, the CTC read channel 0c04 is connected to the VM TCP/IP write channel 0c05. Similarly, CTC write (0c05) is connected to VM TCP/IP read (0c04).

You can write the define and couple commands into the CMS PROFILE EXEC A script. The LINUX for S/390 virtual machine must always be IPLed as CMS before IPLing as LINUX in order for these commands to take effect.

Instead of connecting to the VM TCP/IP user ID, you can connect to any other virtual machine in which a LINUX for S/390, z/OS, OS/390, or VSE system is running.

2. Definitions on the remote side

Set up the TCP/IP on the remote side, as described in the reference manuals. This will vary depending on which operating system is used on the remote side.

Note: It is important that you have IOBUFFERSIZE 65535 defined because the LINUX for S/390 CTC driver works with 64k internally.

3. Activation on the remote side

Activate the channels on the remote side. This again will vary depending on the operating system used on the remote side.

4. Activation on the LINUX for S/390 side

The syntax of this command is:



down deactivates the interface

```
ifconfig ctc0 10.0.51.3 pointopoint 10.0.50.1 netmask 255.0.0.0 mtu 32760
```

```
ifconfig escon0 10.0.0.1 pointopoint 10.0.50.1 netmask 255.0.0.0 mtu 32760
```

Chapter 5. LINUX for S/390 IUCV device driver

The Inter-User Communication Vehicle (IUCV) is a VM/ESA communication facility that enables a program running in one virtual machine to communicate with another virtual machine, or with a control program, or even with itself. The communication takes place over a predefined linkage called a path.

The LINUX for S/390 IUCV device driver is a network device, which uses IUCV to connect LINUX kernels running on different VM user IDs, or to connect a LINUX kernel to another VM guest such as a TCP/IP service machine.


IUCV features

The following features are supported:

- Multiple output paths from a LINUX guest
- Multiple input paths to a LINUX guest
- Simultaneous transmission and reception of multiple messages on the same or different paths
- Network connections via a TCP/IP service machine gateway

IUCV kernel parameter syntax

The driver must be loaded with the IDs of the guest machines you want to connect to:



Parameter:

userid Name of the target VM guest machine

IUCV kernel parameter example

The following picture shows the possible connection of two LINUX for S/390 machines:

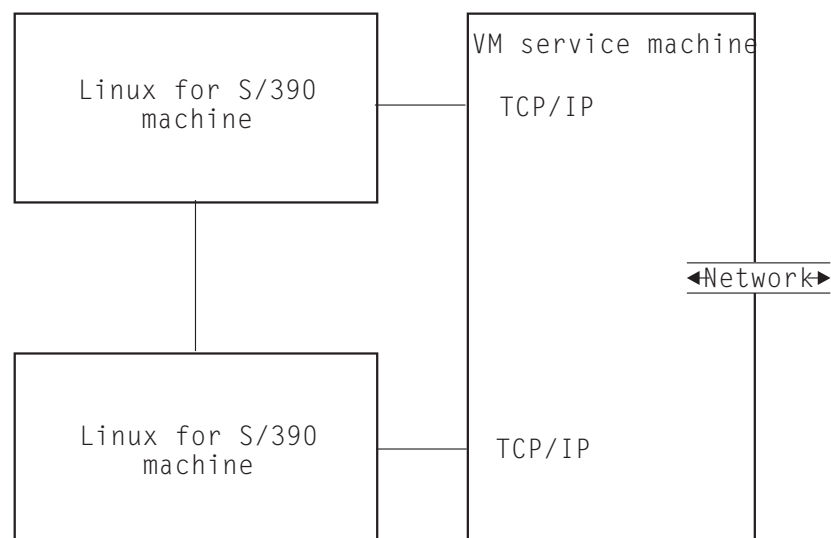


Figure 4. Connection of two systems using IUCV


The command

```
iucv=vmtcpid,linux2
```

connects the LINUX system to the TCP service machine and the other LINUX system.

IUCV module parameter syntax

The driver must be loaded with the IDs of the guest machines you want to connect to:



Parameter:

userid Name of the target VM guest machine

IUCV module parameter example

The example of “IUCV kernel parameter example” on page 26 could be set up by starting the IUCV module with:

```
insmod netiucv iucv=vmtcpid,linux2
```

IUCV – Preparing the connection

This is an additional task that you must perform before you can use the IUCV network link. If LINUX is being used as a network hub instead of VM TCP/IP, the concepts discussed remain the same, though the syntax will be different.

The following steps must be undertaken in VM:

1. Obtain a subnet from your TCP/IP communications staff. It is important that the subnet used by your LINUX guests not be the same as that used by VM on the LAN. It is a separate network and should be treated as such.
2. Take one address from that subnet and assign it to VM. Update your PROFILE TCPIP file with a home entry, device, link, and start statements for each guest, for example:

```
Home
  vm_ip_address link_name1
  vm_ip_address link_name2

Device device_name1 IUCV 0 0 linux_virtual_machine1 A
Link link_name1 IUCV 0 device_name1

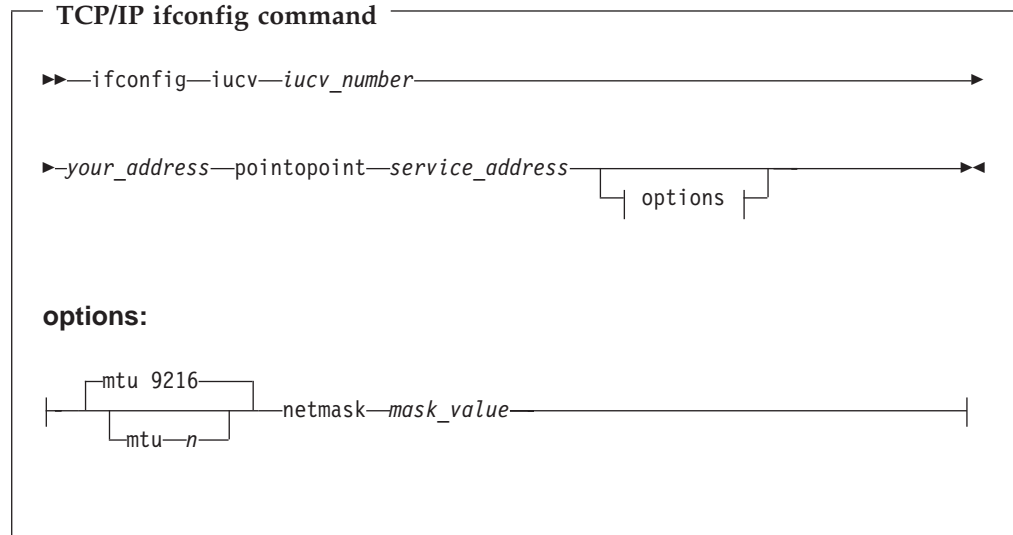
Device device_name2 IUCV 0 0 linux_virtual_machine2 A
Link link_name2 IUCV 0 device_name2

Start device_name1
Start device_name2
```

3. Add the necessary VM TCP/IP routing statements (BsdRoutingParms or Gateway). Use an MTU size of 9216 and a point-to-point host route (subnet mask 255.255.255.255). If you use dynamic routing, but do not wish to run routed or gated on LINUX, update the VM ETC GATEWAYS file to include “permanent” host entries for each LINUX guest.
4. Bring these updates on-line by using OBEYFILE or by recycling TCPIP and/or ROUTED as needed.
5. Add the statement
IUCV ALLOW

to your VM user directory entry.

The LINUX commands needed to start communications through a TCP/IP service machine are:



Parameters:

iucv_number

Path number (for example 0)

your_address

TCP/IP address of your machine

pointopoint

required to establish a point-to-point connection to a service machine

service_address


Address of the TCP/IP service machine to connect to

n

maximum transfer unit size. The default is 9216, which is suitable for use with the S/390 Virtual Image Facility for LINUX (VIF). The maximum value is 32764.

mask_value

Mask to identify addresses served by this connection



Parameters:

iucv_number

Path number defined above

inetd command

```
(1)
▶▶—inetd—————▶▶
```

Notes:

- 1 Not required if the IUCV driver is started during boot.

The commands needed to start direct communications to another guest are:

user-to-user ifconfig command

```
▶▶—ifconfig—iucv—iucv_number—————▶
▶—guest_0_address—pointopoint—guest_1_address—————▶▶
```

Parameters:

iucv_number

Path number (for example 0)

guest_0_address

TCP/IP address of your machine

guest_1_address

TCP/IP address of target machine

IUCV – Further information

The standard definitions in the VM TCP/IP configuration files apply.

For more information of the VM TCP/IP configuration see: *VM/ESA TCP/IP Planning and Customization* , SC24-5847-01.

IUCV restrictions

- This device driver is only available to LINUX for S/390 systems running as guests under VM/ESA.

Chapter 6. LINUX for S/390 Console device drivers

The S/390 hardware requires a line-mode terminal (the hardware console) for overall system control. The LINUX for S/390 console device drivers enable LINUX to use this console for basic LINUX control as well.

You can use a 3215 console instead of the hardware console if LINUX is running under VM or on a P/390.

Both drivers can be compiled into the LINUX kernel. If both are present the appropriate console driver will be chosen at run-time according to the environment:

- In VM, the 3215 console will be enabled and the hardware console suppressed.
- In an LPAR or native environment, the hardware console will be enabled and the 3215 suppressed.

The intended use of the console device drivers is solely to launch LINUX . When LINUX is running, the user should access LINUX for S/390 via Telnet, because the console is a line-mode terminal and is unable to support applications such as vi. Therefore it is strongly recommended that you assign *dump* to the *TERM* environment variable so that at least applications like less give appropriate output.

The hardware console driver code is separated into three different files (with header files) - one for the low level hardware console handling, another one to make use of the low level driver as a LINUX console and the third as a bridge between the generic terminal code of the kernel and the low level hardware console driver. Note that there are different options that must be selected during kernel configuration to enable the LINUX terminal on the hardware console or to enable the LINUX console on the 3215 console.

The main feature of the hardware console driver is the ability to support the kernel and user space in parallel. Accordingly there are two different interfaces within the kernel: console and TTY (terminal) to interact with the hardware console driver. The TTY processes all output requests from outside of the kernel - user space - the consoles write only kernel messages, but are not able to get input from the operator. The hardware console device driver uses a small subset of the hardware console API to interact with a command interface on the SA/SE (stand alone support element) or a connected HMC (hardware management console).

Console features

- Provides a line mode typewriter terminal.
- Console output on the first 3215 terminal (VM only).


Console kernel parameter syntax

The hardware console device driver does not require any parameters.

The 3215 console device driver does not require any parameters if it is used under VM. If it is used with a P/390 system, you have to specify the *condev* kernel parameter. This supplies the device driver with the subchannel number of the 3215

device. The reason that this parameter is needed is that there is no guaranteed method of recognizing a 3215 device attached to a P/390.

The kernel parameter syntax is:



where *cuu* is the device 'Control Unit and Unit' address, and may be expressed in hexadecimal form (preceded by 0x) or in decimal form.

Note: Early releases of the driver will not accept this parameter in hexadecimal form.

Console kernel examples

```
condev=0x001f
```

or

```
condev=31
```

Both of these formats tell the device driver to use device number hex 1F for the 3215 terminal.

Using the console

The console is a line mode terminal. The user enters a complete line and presses enter to let the system know that a line has been completed. The device driver then issues a read to get the completed line, adds a new line and hands over the input to the generic TTY routines.

Console special characters

The console does not have a control key. That makes it impossible to enter control characters directly. To be able to enter at least some of the more important control characters, the character '^' has a special meaning in four cases:

- The two character input line ^c is interpreted as a Ctrl+C. This is used to cancel the process that is currently running in the foreground of the terminal.
- The two character input line ^d is interpreted as a Ctrl+D. This is used to generate an end of file (EOF) indication.
- The two character input line ^z is interpreted as a Ctrl+Z. This is used to stop a process.
- The two characters ^n at the end of an input line suppresses the automatic generation of a new line. This makes it possible to enter single characters, for example those characters that are needed for yes/no answers in the ext2 file system utilities.

If you are running under VM without a 3215 console you will have to use the CP VINPUD command to simulate the **ENTER** and **SPACE** keys.

The **ENTER** key is simulated by entering:

```
#CP VInput VMSG \n
```

The **SPACE** key is simulated by entering:

```
#CP VInput VMSG \n      (two blanks followed by \n).
```

If the special characters do not appear to work, make sure you have the correct codepage in your terminal emulator. One known to work is codepage 037 (United States).

Console 3270 emulation

If you are accessing the VM console using the **x3270** emulator, then you should add the following settings to the `.XDefaults` file in order to get the correct code translation:

```
! X3270 keymap and charset settings for Linux
x3270.charset: us-intl
x3270.keymap: circumfix
x3270.keymap.circumfix: :<key>asciicircum: Key("^")\n
```

Console – Use of VInput

Note: 'VInput' is a VM CP command. It may be abbreviated to 'VI' but should not be confused with the LINUX command 'vi'.

If you use the hardware console driver running under VM it is important to consider how the input is handled. Instead of writing into the suitable field within the graphical user interface at the service element you have to use the VInput command provided by VM. The following examples are written at the input line of a 3270-Terminal or 3270-Terminal-Emulation (for example, x3270).

Note that, in the examples, capitals within VInput and its parameters processed by VM/CP indicate the characters you have to type. The lower case letters are optional and are shown for readability.

```
VInput VMSG LS -L
```

(the bash will call `ls -l` after this command was sent via VInput to the hardware console as a non-priority command - VMSG).

```
VInput PVMG ECHO *
```

(the bash will execute `echo *` after this command was sent via VInput to the hardware console as a priority command - PVMG).

The hardware console driver is capable to accept both if supported by the hardware console within the specific machine or virtual machine. Please inspect your boot messages to check the hardware console's capability of coping with non-priority or priority commands respectively on your specific machine. Remember that the hardware console is unable to make its own messages available via `dmesg`.

Features of VInput (found out experimentally).

1. Use `''''` to output a single `''`.

```
VInput example: VInput PVMG echo ""Hello world, here is ""$0
```

(on other systems: `echo "Hello world, here is "$0`)

2. Do not use # within VInput commands.
is interpreted as a control character by VM CP. This character and all following characters are read by CP and are not transmitted to the hardware console or its driver.
3. All characters in lower case are converted by VM to upper case.
If you type VInput VMSG echo \$PATH, the driver will get ECHO \$PATH and will convert it into echo \$path. LINUX and the bash are case sensitive and cannot execute such a command. To resolve this problem, the hardware console uses an escape character (%) under VM to distinguish between upper and lower case characters. This behavior and the escape character (%) are adjustable at build-time by editing the driver sources, or at run-time by use of the ioctl interface. Some examples:
 - input: VInput VMSG ECHO \$PATH
output: echo \$path
 - input: vi vmsg echo \$%PATH%
output: echo \$PATH
 - input: #cp vi pvmsg echo ""%7/27/00ello, here is ""\$0 #name of this process
output: CP VI PVMSG ECHO ""%7/27/00ELLO, HERE IS "\$0
NAME OF THIS PROCESS
HPCMD001E Unknown CP command: NAME
echo "%Hello, here is "\$0
%Hello, here is -bash

Console limitations

- The 3215 driver only works in combination with VM. In a single image or in LPAR mode the 3215 terminal device driver initialization function just exits without registering the driver.
- Due to a problem with the translation of code pages (500, 037) on the host, the pipe command character (|) cannot be intercepted by the console. If you need to use this command execute it from a Telnet session.
- Displaying large files might cause some missing sections within the output because of the latency of the hardware interface employed by the device.
- In native or LPAR environments, you occasionally have to use the **Delete** button of the GUI on the Service Element or Hardware Management Console to enable further output. This is relevant to:
 - SE version 1.6.1 or older on G5, G6, and Multiprise 3000.
 - SE version 1.5.2 or older on G3, G4, and Multiprise 2000.
- Messages concerning the hardware console operation generated by the hardware console driver cannot be provided to the syslog and are therefore unavailable with dmesg.
- Output from the head/top is deleted if the amount exceeds approximately 30 Kbyte per LPAR (or image) on SE or HMC.
- Applications such as vi are not supported because of the console's line-mode nature.

Part 3. LINUX for S/390 Object-Code-Only device drivers

These chapters describe the device drivers available for optional S/390 hardware.

The drivers described are:

- “Chapter 7. LINUX for S/390 LCS Device Driver” on page 37
- “Chapter 8. LINUX for S/390 Gigabit Ethernet device driver” on page 43

Licence conditions

These drivers are subject to licence conditions as reflected in: “International License Agreement for Non-Warranted Programs” on page 115.

Chapter 7. LINUX for S/390 LCS Device Driver

This driver is subject to licence conditions as reflected in: "International License Agreement for Non-Warranted Programs" on page 115.

This LINUX network driver supports LAN Channel Station (LCS) Ethernet and Token Ring access through the OSA-2™ card.

The LCS network interface has two channels, one read channel and one write channel. This is very similar to the S/390 CTC interface. The read channel must have an even cuu number and model 0x3088. The write channel has a cuu number one greater than the read cuu number, and also has a model of 0x3088. Only certain cuu types are supported so as not to clash with a CTC control unit type.

The driver always has a read outstanding on the read subchannel. This is used to receive command replies and network packets (these are differentiated by checking the type field in the LCS header structure). Any network packets that arrive during the start-up and shutdown sequence have to be discarded. During normal network I/O, the driver will intermittently retry reads in order to permanently keep a read outstanding on the read channel. (This is in case an -EBUSY or similar occurs, in which case the driver would stop receiving network packets.)

The default configuration is to use software statistics, with IP checksumming off (this improves performance) and to have network hardware checking using a CRC32 check (CRC64 for FDDI) which should guarantee integrity for normal use. However, financial institutions or similar might want the additional security of IP checksumming.

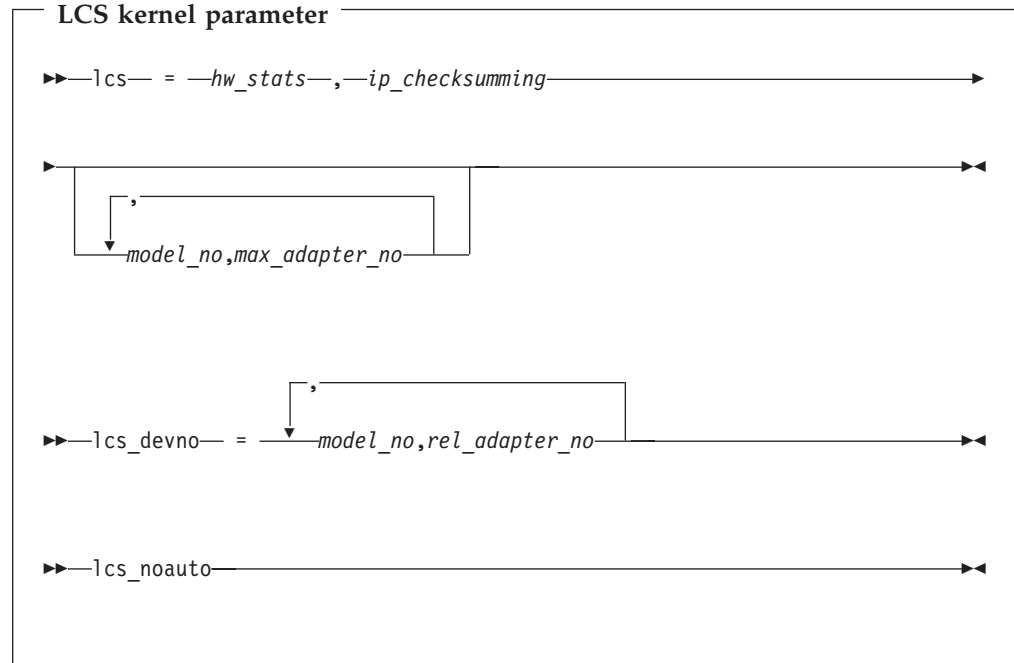
Additional CUU model types can be added later so that new LCS compatible cards will be supported even if not available when the driver was developed.

LCS features

- Supports Ethernet and Token Ring
- Auto detects whether card is connected to Token Ring or Ethernet
- Configurable to compile as a module or directly into the kernel
- Can be configured from the kernel parameter line or via insmod parameters (if a module).

LCS kernel parameter syntax

If the LCS driver is compiled directly into the kernel, the LCS boot parameters are as follows:



The meanings of these parameters are:

hw_stats

If set to 1 it is equivalent to the **use_hw_stats** keyword for the module call.
If set to 0 it is ignored.

ip_checksumming

If set to 1 it is equivalent to the **do_sw_ip_checksumming** keyword for the module call. If set to 0 it is ignored.

model_no,max_adapter_no

is identical to the **additional_model_info** keyword described for modules.

model_no,rel_adapter_no

is identical to the **devno_portno_pairs** keyword described for modules.
The keyword **lcs_devno** is short because of limitations in the allowable length of kernel parameter names.

lcs_noauto

Put this parameter in the kernel parameter line if you want to set auto-detection off.

It is important that the parameters are entered in pairs (2, 4, 6 or 8 parameters) as the cu model and max rel adapter no must go together.

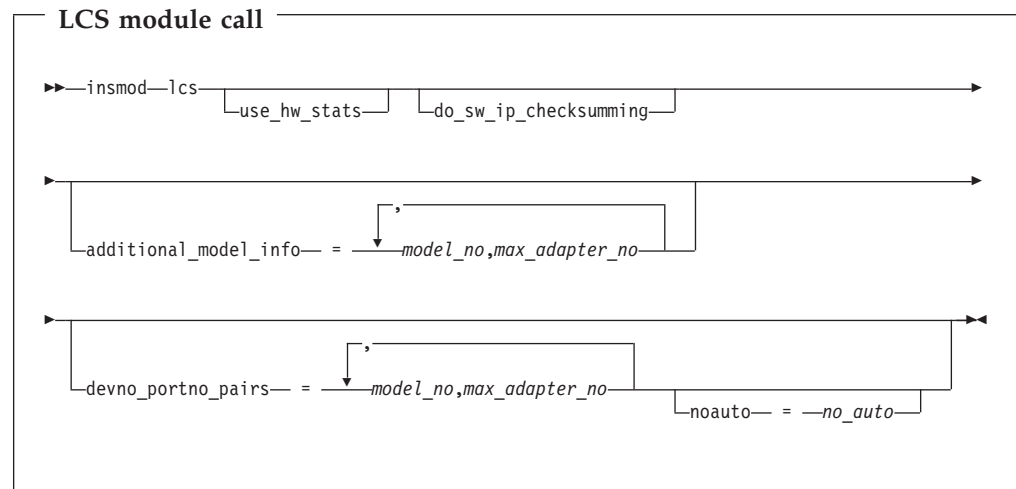
LCS kernel parameter example

```
lcs 1,1,97,8
```

This kernel parameter will collect network statistics from the hardware, set IP checksumming on, define the model number as 0x61 (97 decimal), and make an attempt to detect 8 ports on that model.

LCS module parameter syntax

The following are the LCS device driver module parameters:



The meanings of these keywords are:

use_hw_stats

Gets network statistics from LANSTAT LCS primitive as opposed to doing it in software. This is not recommended as it is incompatible with many network drivers

do_sw_ip_checksumming

Performs IP checksumming on inbound packets in the software. Normally not required because Ethernet CRC32 is usually more than adequate (except perhaps for financial institutions).

additional_model_info

Made up of sets of model/maximum relative adapter number pairs (see example below).

devno_portno_pairs

Takes `devno,rel_adapter_no(port)` pairs. Relative adapter number of -1 indicates that you should not use this adapter. This can be used to force certain CHPIDs to use a particular port number if the LCS protocol returns an incorrect one.

noauto

Set `noauto=1` if you want to set auto-detection to off. You then have to explicitly configure LCS devices with the `devno_portno_pairs` module parameter.

LCS module parameter example

```
insmod lcs additional_model_info=0x70,3,0x71,5
devno_portno_pairs=0x1c00,0,0x1c02,1,0x1d00,-1
```

tells the LCS device driver to:

- look for 3 ports on a 3088 model 70 and 5 ports on a model 71
- only use port 0 if available for the device numbers 0x1c00 and 0x1c01
- only use port 1 if available for the device numbers 0x1c02 and 0x1c03
- do not under any circumstances use the device at 0x1d00 and 0x1d01 as an LCS device.

LCS restrictions

- LINUX for S/390 cannot run with a root file system mounted via NFS when the network connectivity is established via LCS, because the LCS driver is delivered as an object-code-only module.
- To use OSA-2 devices when running LINUX for S/390 on a basic mode machine (no LPARs) you need to specify an `ipldelay=xyz` boot parameter. We recommend a value between 2m and 5m for xyz for the OSA-2 card to settle down after LOAD.
- Currently, there is only support for up to 16 Token Ring or Ethernet devices. However, we strongly recommend that you do not share devices with production systems.

LCS limitations

- FDDI is untested and the code shipped with kernel patch 2.2.15 is unlikely to work on FDDI
- Because LCS does not appear to tell the driver that a port is busy, it is sometimes necessary to force the driver not to use the ports. The `devno_portno_pairs` or `lcs_devno` kernel parameters are used to do this.
- The most problematic area for the code is starting up and shutting down the driver.

This is primarily due to the fact that network packets can be received during the start-up process before receiving the `lanstat` command to get the mac address. This can happen earlier if the card wasn't previously shut down properly.

If the card is being very troublesome, use `ifconfig` to switch it on and off.


If this fails, compile the driver as a module. Use `insmod` and `rmmod`, as these are guaranteed to call the start-up and shutdown routines, whereas the kernel keeps a reference count (doing `ifconfig` up twice will call the start-up routine only once).

LCS – Common set up problem

The LCS device driver sometimes has a set up problem. The same port on the OSA/2 card is taken twice by the LCS driver. This port can be communicated to via two or more pairs of device numbers. The driver attempts to determine the port number from the low byte of the device number, however the LCS microcode does not indicate that the port is already in use. If the first attempt is wrong (port already in use), the driver may use the same port twice with different pairs of device numbers. This can be better explained via the following diagrams.


Normal case:

Device Numbers




Abnormal case - part 1:

Device Numbers



Abnormal case - part 2 (two or more connections on the same port):

Device Numbers



The common symptoms of this problem are:

1. Duplicate mac addresses (most common)
2. Two or more network interfaces detected of the same type despite the fact that port X is configured as Token Ring and port Y as Ethernet
3. Disconnecting network cable from one port disables network traffic on two or more ports.

If this problem occurs, use the `devno_portno_pairs` parameter to force the problematical device numbers to specific ports.

Chapter 8. LINUX for S/390 Gigabit Ethernet device driver

This driver is subject to licence conditions as reflected in: "International License Agreement for Non-Warranted Programs" on page 115.

This LINUX network driver supports the OSA-Express Gigabit Ethernet card. This enables the S/390 to connect to Gigabit Ethernet networks, or to Ethernet or Fast Ethernet networks through a switch. ³

GbE features

The following features are supported:

- Autosensing of devices
- Primary and secondary routers
- Priority queueing
- Individual device configuration.

GbE module parameter syntax

Before the Gigabit Ethernet device driver can be loaded the QDIO protocol driver must be loaded.

The command to load the protocol driver is:

qdio module call

►► insmod qdio ◀◀

When the QDIO protocol has been loaded the Gigabit Ethernet driver (qeth) is loaded with the command:


qeth module call using default options

►► insmod qeth ◀◀

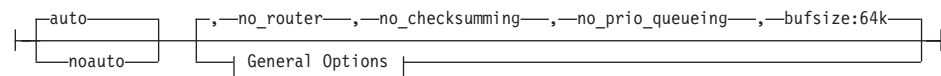
or:

3. Please read the licence on page "International License Agreement for Non-Warranted Programs" on page 115.

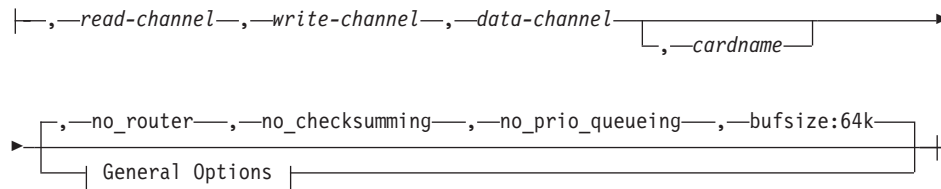
qeth module call specifying options



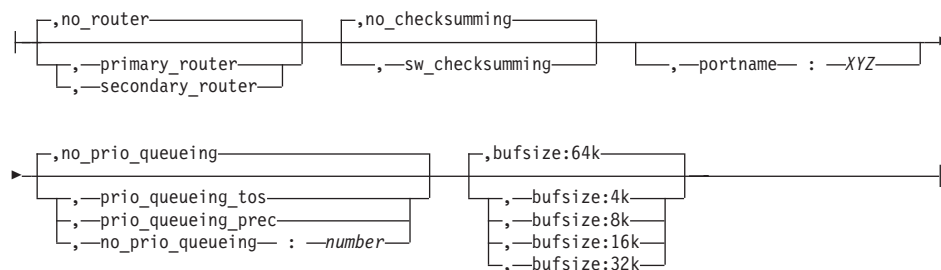
Driver Options:



Card Options:



General Options:



Notes:

- 1 All options except the first used must be preceded by a comma.

Note: All characters must be entered in lower case as shown, except in hexadecimal numbers where either case may be used.

The meaning of the parameters of this command is as follows:

auto | noauto

Specifies whether the driver is to search for all subchannels (auto), or is only to use device numbers specified in *card options*

primary_router | secondary_router | no_router

Specifies whether the device is used to interconnect networks. A "Primary router" is the principal connection between two networks; a "Secondary router" is used as backup in case of problems with the primary. Both of these options require the LINUX system to be configured as a router. The default for this parameter is "No router" – the Gigabit Ethernet card will only be used to connect the LINUX for S/390 system to a single network.

sw_checksumming | no_checksumming

Specifies whether error detection is to be performed by the driver, or is not required.

prio_queueing_tos | prio_queueing_prec | no_prio_queueing |

no_prio_queueing: *number*

Specifies the type of priority queuing to be used. See "GbE queuing" on page 48 for details. **no_prio_queueing** is equivalent to **no_prio_queueing:2**

bufsize:*numberk*

Specifies the size of the qdio inbound buffer. Valid values for *number* are 4, 8, 16, 32 and 64. The default is 64.

portname: XYZ

Identification of the port for sharing by other OS images, for example the 'Portname' dataset used by OS/390. XYZ is up to 8 characters long.

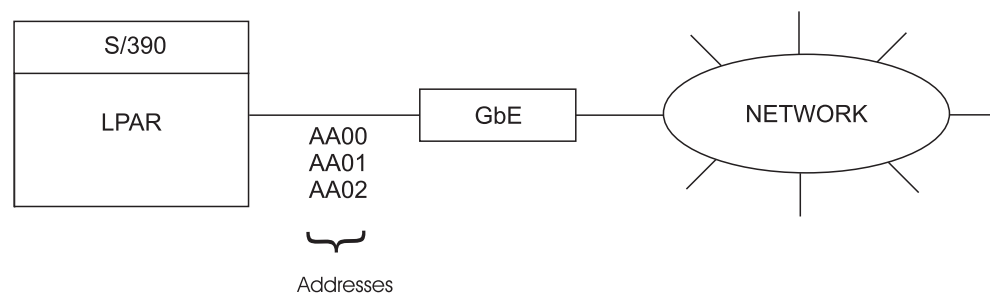
<card options> are used to override the global options for a particular device. These are also comma-separated lists. The card is identified by its three device numbers (in decimal, or in hexadecimal prefixed with 0x). These numbers may be followed by a device name, and may be followed by any of the **<driver options>** keywords except **auto** or **noauto**.

GbE examples

1: Basic configuration

In this example a single Gigabit Ethernet card is being used to connect a LINUX for S/390 system to a network.

Hardware configuration – Gigabit Ethernet connecting LINUX for S/390 to a network.



Software configuration – Gigabit Ethernet connecting LINUX for S/390 to a network.

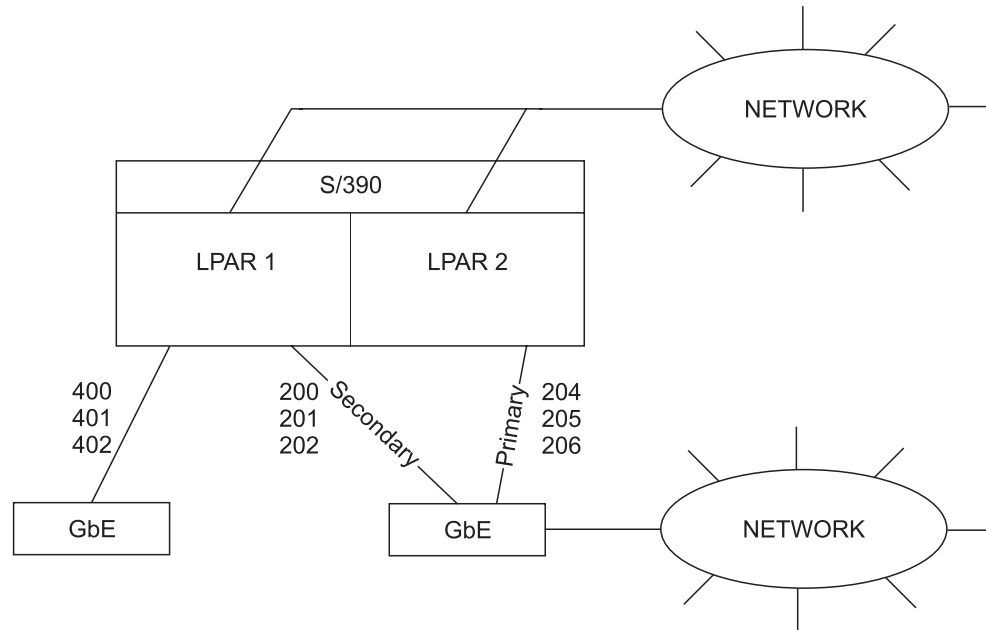
The load commands for this configuration are:

```
insmod qdio
insmod qeth qeth_options=noauto,0xAA00,0xAA01,0xAA02
```

2: Router configuration

This example shows how LINUX systems running on different LPARs in an S/390 may use Gigabit Ethernet cards to communicate with a network or to act as a router between networks.

Hardware configuration – Gigabit Ethernet and LINUX for S/390 as router.



In this example it is assumed that LINUX is configured as a router in both LPARs.

Software configuration – Gigabit Ethernet and LINUX for S/390 as router.

LPAR 1 – uses the first subchannel group as a network client, and the second subchannel group as a backup router for LPAR 2:

```
insmod qdio
insmod qeth qeth_options=noauto,0x400,0x401,0x402,0x200,0x201,0x202,secondary_router
```

LPAR 2 – uses the third subchannel group as a primary router:

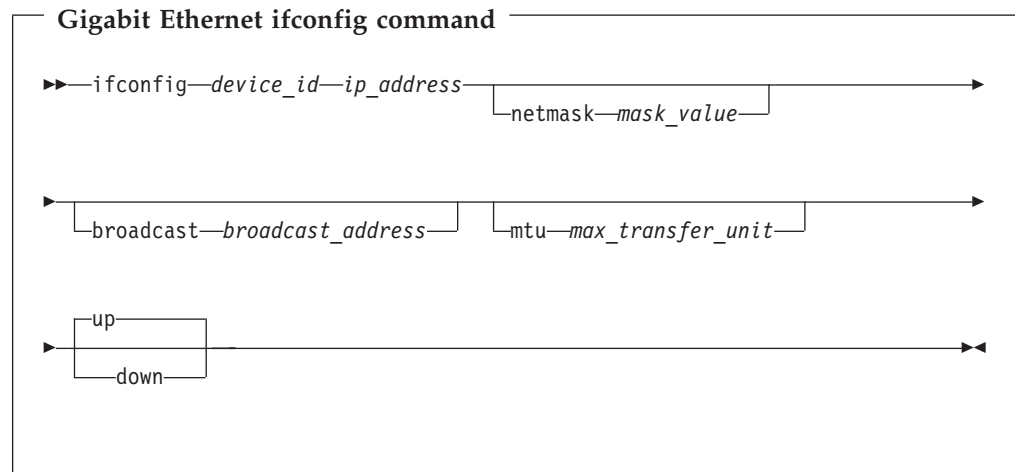
```
insmod qdio
insmod qeth qeth_options=0x204,0x205,0x206,primary_router
```

GbE – Preparing the connection

Activating the Gigabit Ethernet connection

The network devices can be activated with the `ifconfig` command. It is necessary to define the right MTU size for the channel device, otherwise it will not work properly. You must use the same MTU size (default 1492) that is defined on the remote side:

The syntax of this command is:



Where:

device_id
identifies the interface

ip_address
is the IP address of the remote side.

mask_value
is the IP network mask for this interface

broadcast_address
is the address used to send to all devices on the connection

max_transfer_unit
is the size of the largest block which may be carried

up activates the interface

down deactivates the interface

An example of the use of ifconfig is:

```
ifconfig eth0 192.168.100.11 netmask 255.255.255.0 broadcast 192.168.100.255 mtu 1492 up
```

or, more simply:

```
ifconfig eth0 192.168.100.11
```

GbE device recognition

With autosensing on any device addresses specified in the driver parameter are checked, followed by a scan of all other device subchannels in ascending order. With autosensing switched off only the device addresses specified in the driver parameter are checked. This may be necessary to prevent the card from taking other device numbers if the specified numbers cannot be used.

A subchannel is checked to see if:

- it is a **QDIO** eligible subchannel
- the channel is not already in use
- it is on an unused CHPID (unless explicitly specified in options)
- the device type and model are known

When one subchannel has been verified a search for two more subchannels is carried out on the same CHPID. The same criteria apply. When a group of three subchannels has been found, the driver signals a request for these and the card is set up.

GbE restrictions

- Currently the I/O Layer does not provide CHPID information, so all devices are seen as on one CHPID. This means only one device will automatically be detected.
- The MTU range is 576 – 56000. Standard sizes used are 576, 1492, 1500, 8992 and 9000.
- The QDIO based Ethernet Driver `qeth.o` does not work with full autosensing in any situation. It uses different heuristics for finding subchannels when fully autosensing compared to the case where device numbers are specified. In the latter case, the specified device numbers are always used and the `qeth` driver recognizes the device. When fully autosensing, the driver is sometimes not able to recognize a group of three subchannels, as it does not try all permutations of the subchannels. The problem will appear when:
 1. there are only three subchannels
 2. the device numbers of these are consecutive and start with an odd number, for example 0x1103, 0x1104 and 0x1105

In this case full autosensing does not work, but

```
insmod qeth qeth_options=0x1103,0x1104,0x1105
```

will work. To avoid this problem either attach four subchannels to each VM guest or LPAR using the card, or ensure that the subchannels start with an even number.

GbE queuing

The Gigabit Ethernet card has four output queues (queues 0 to 3) in central storage. The card gives these queues different priorities (queue 0 having the highest priority) which is relevant mainly to high traffic situations. When there is little traffic queuing has no impact on processing.

The device driver can put data on one or more of the queues. By default it uses queue 2 for all data. However, the driver can look at outgoing IP packets and select a queue for the data according to the IP type of service or IP precedence fields. (These fields are part of the IP datagram header.)

Some applications use these fields to tag data. The mapping the driver performs between IP type of service is as follows:

IP type of service	queue used when IP TOS queueing is switched on
not important	3
low latency	0
high throughput	1
high reliability	2
everything else	2

When IP precedence was selected as queueing type, the two most significant bits of each IP header precedence field are used to determine the queue for this packet.


GbE background – QDIO

The QDIO protocol governs the interface between the S/390 and the Gigabit Ethernet card.

The Gigabit Ethernet is optimized for low latency and SMP environments (rather than high throughput).

For Gigabit Ethernet devices three I/O channels must be available to the driver. One channel is for control reads, one for control writes, and the third is for data.

Example card layout



Part 4. Installation commands and parameters

Chapter 9. Useful LINUX commands

This chapter describes:

- Commands that are used during installation and initial start-up of LINUX for S/390
- Commands created for LINUX for S/390

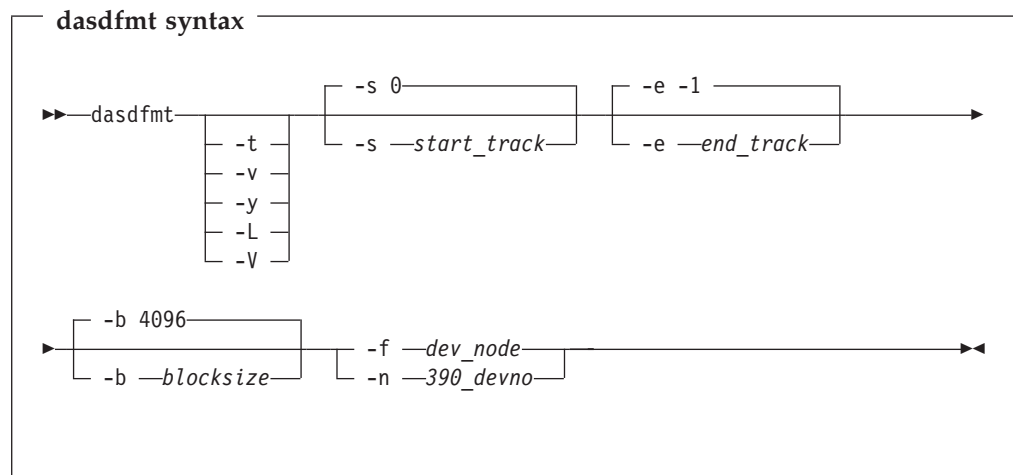
dasdfmt - Format a DASD

Usage

This tool is used to low-level format direct access storage devices (DASD). Note that `dasdfmt` is only able to format DASDs that have already been formatted using another disk formatting utility. If you have unformatted DASD in your system, use a device support facility such as ICKDSF to initially format the DASD.

`dasdfmt` uses an `ioctl` call to the DASD driver to format tracks. A start and end track for formatting can be specified, as well as a blocksize (hard sector size). Remember that the formatting process can take quite a long time.

Format



`dasdfmt -help` prints out an overview of the syntax.

The parameters are:

-f *dev_node*

Specifies the device node in the file system. This must be the whole device, not a partition.

-n *390_devno*

Specifies the device number of the disk to format.

The following parameters are necessary, however if you do not specify their values you are prompted for them. You can use the default values by pressing the <enter> key:

-s *start_track*

Specifies the start track of the formatting. A value of 0 (first track of disk) is the default value.

-e *end_track*

Specifies the last track of the formatting. A value of -1 means the last track on the disk and is the default value.

-b *block_size*

Specifies the blocksize. The minimum blocksize is 512 bytes and increases in powers of 2 (512, 1024, 2048, 4096 and so on). The default blocksize is

4096. We recommend setting the blocksize to 1024 or higher (ideally 4096) because the ext2fs file system uses 1KB blocks and 50% of the capacity will be unusable if the DASD blocksize is only 512 bytes.

The following parameters are optional:

- v Prints out more messages.
- y Omits the security prompt and formats the disk directly (for batch use by daring people!).
- t Switches to a test mode, the DASD will not be formatted.
- L Don't write disk label (default is write label LNX1).
- V Print version of dasdfmt.

Examples

To format two whole disks with a blocksize of 4096, one the device nodes /dev/dasda and the other at address 0x0293:

```
dasdfmt -b 4096 -f /dev/dasda
dasdfmt -b 4096 -n 0293
```

ifconfig - Configure a network interface

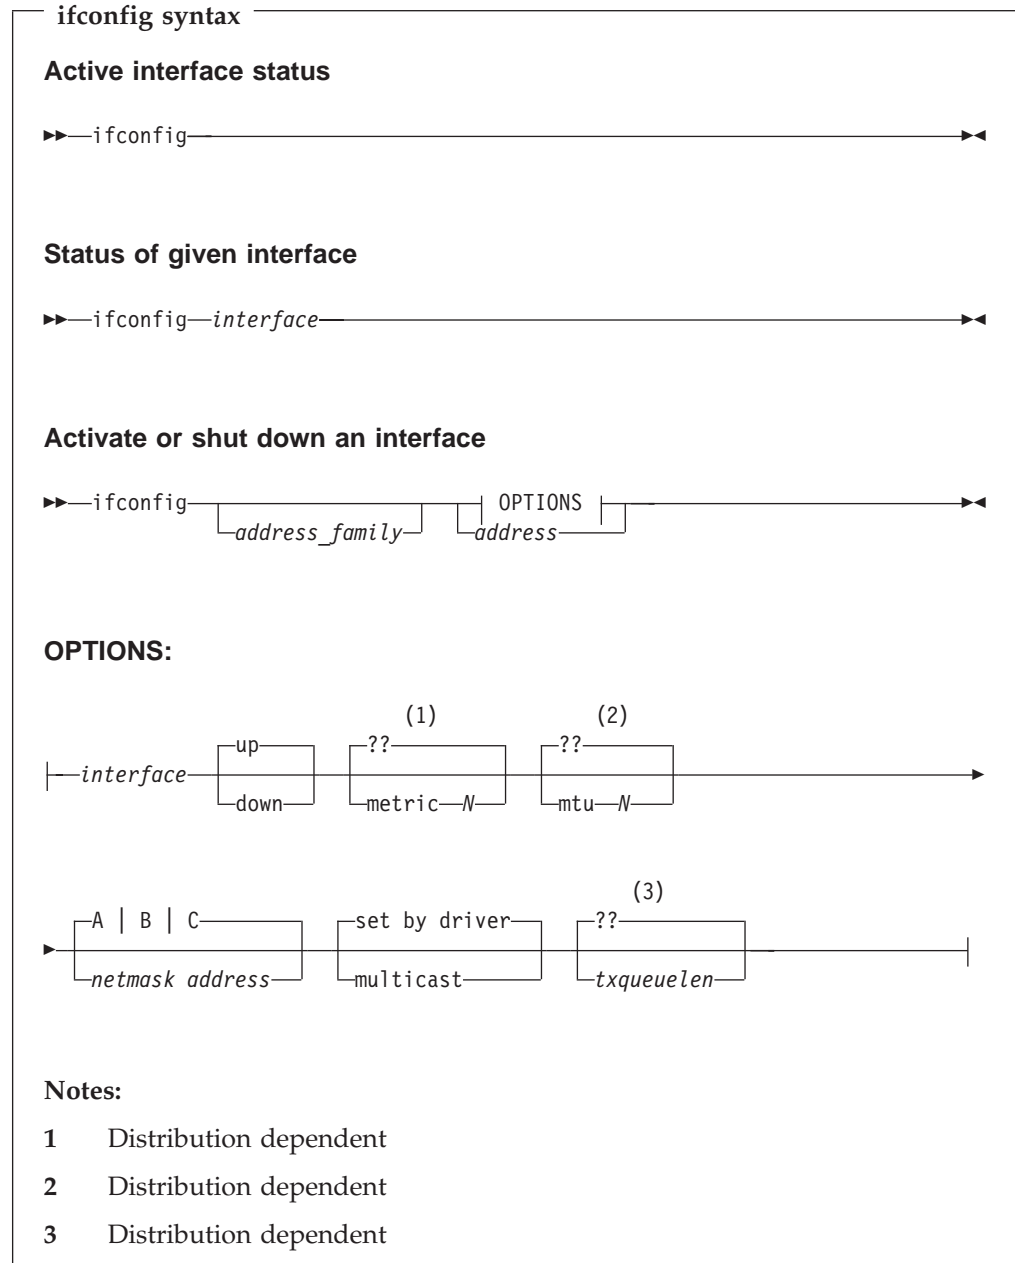
Usage

`ifconfig` is used to configure the kernel-resident network interfaces. It is used at start-up time to set up interfaces as necessary. After that, it is usually only needed when debugging or when system tuning is needed.

- If no arguments are given, `ifconfig` displays the status of the currently active interfaces.
- If a single interface argument is given, it displays the status of the given interface only
- Otherwise, it configures an interface. The configurable interfaces for S/390 are:
 - `iucv`
 - `ctci` ($i = 0$ to 7)
 - `esconj` ($j = 0$ to 7)

Note: Since kernel release 2.2 there are no explicit interface statistics for alias interfaces anymore. The statistics printed for the original address are shared with all alias addresses on the same device. If you want per-address statistics you should add explicit accounting rules for the address using the `ipchains(8)` command.

Format



address_family

If the first argument after the interface name is recognized as the name of a supported address family, that address family is used for decoding and displaying all protocol addresses.

On S/390, supported address families include:

- inet

interface

The name of the interface. This is usually a driver name followed by a unit number, for example eth0 for the first Ethernet interface.

On S/390, supported interfaces include:

- `escon0` - `escon7`
- `ctc0` - `ctc7`
- `iucv0`
- `lo0` (loopback device)
- `eth0`
- `tr0`

up This flag causes the interface to be activated. It is implicitly specified if an address is assigned to the interface.

down This flag causes the driver for this interface to be shut down.

metric N

This parameter sets the interface metric.⁴

mtu N This parameter sets the maximum transfer unit (MTU) of an interface.

netmask addr

Set the IP network mask for this interface. This value defaults to the usual class A, B or C network mask (as derived from the interface IP address), but it can be set to any value.

multicast

Set the multicast flag on the interface. This should not normally be needed as the drivers set the flag correctly themselves.

address

The IP address to be assigned to this interface.

txqueuelen length

Set the length of the transmit queue of the device. It is useful to set this to small values for slower devices with a high latency (modem links, ISDN) to prevent fast bulk transfers from disturbing interactive traffic like telnet too much.

Files:

`/proc/net/socket`
`/proc/net/dev`

Examples

To start or modify an ESCON interface in LINUX:

```
ifconfig escon0 x.x.x.x pointopoint y.y.y.y netmask 255.255.255.255 mtu mmmmm
```

where:

4.

Metric: Cost of an OSPF interface. The cost is a routing metric that is used in the OSPF link-state calculation. To set the cost of routes exported into OSPF, configure the appropriate routing policy.

- Range: 1 through 65,535
- Default: 1

All OSPF interfaces have a cost, which is a routing metric that is used in the OSPF link-state calculation. Routes with lower total path metrics are preferred over those with higher path metrics. When several equal-cost routes to a destination exist, traffic is distributed equally among them. The cost of a route is described by a single dimensionless metric that is determined using the following formula:

$$\text{cost} = \text{ref-bandwidth} / \text{bandwidth}$$

Where ref-bandwidth is the reference bandwidth. Its default value is 100 Mbps (which you specify as 100000000), which gives a metric of 1 for any bandwidth that is 100 Mbps or greater.

x.x.x.x is the IP address of the LINUX side

y.y.y.y is the IP address of the remote partner OS/390

mmmmm

is the MTU size which could be up to 32760 - make sure the other side of the channel uses the same MTU size

The ESCON CTC device addresses are defined in the kernel parameter file, or when loading the module:

```
..... ctc=0,0xdd,0xxyy,escon0
```

Another example, showing how to set up an ethernet connection:

```
ifconfig eth0 192.168.100.11 netmask 255.255.255.0 broadcast 192.168.100.255 mtu 1492 up
```

or, simply:

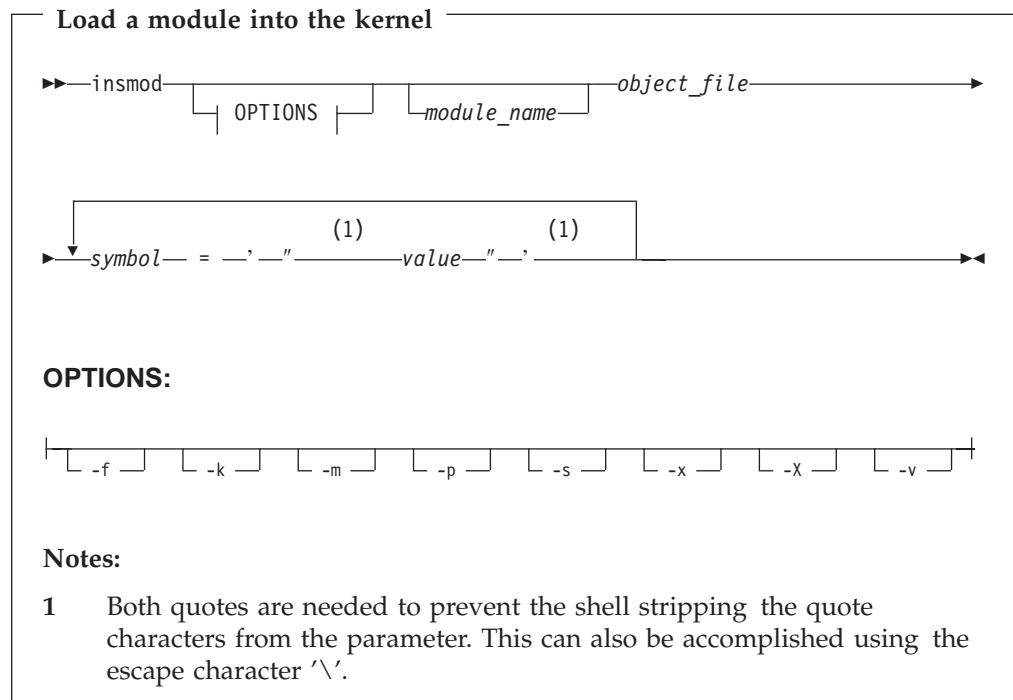
```
ifconfig eth0 192.168.100.11
```

insmod - Load a module into the LINUX kernel

Usage

insmod installs a loadable module in the running kernel. It tries to link a module into the kernel by resolving global symbols in the module with values from the kernel's symbol table. If the object file name is given without extension, insmod will search for the module in common default directories. The environment variable MODPATH can be used to override this default.

Format



object_file

Name of module source file. (Name by which module is invoked.)

module_name

Explicit name of module if not derived from the name of the source file.

symbol

Name of parameter specific to module.

value

Value of parameter to be passed to module.

-f

Attempt to load the module even if the version of the running kernel and the version of the kernel for which the module was compiled do not match.

-k

Set the auto-clean flag on the module. This flag will be used to remove modules that have not been used in some period of time (usually one minute).

-m

Output a load map, making it easier to debug the module in the event of a kernel panic.

- p** Probe the module to see if it could be successfully loaded. This includes locating the object file in the module path, checking version numbers, and resolving symbols.
- s** Output everything to syslog instead of the terminal.
- X** Export the external symbols of the module.
- x** Do not export the external symbols of the module.
- v** Verbose mode.

Examples

DASD module

```
insmod dasd_mod dasd=192-194,5a10
```

XPRAM module

```
insmod xpram devs=4 sizes=2097152,8388608,4194304,2097152
```

CTC module

```
insmod ctc setup=\"ctc=0,0x0600,0x0601,ctc0\"
```

LCS module

```
insmod lcs additional_model_info=0x70,3,0x71,5
          devno_portno_pairs=0x1c00,0,0x1c02,1,0x1d00,-1
```

QDIO module

```
insmod qdio
```

Gigabit Ethernet module


```
insmod qeth qeth_options=noauto,0x400,0x401,0x402,0x200,0x201,0x202,secondary_router
```

mke2fs - Create a file system

Usage

This utility creates an ext2 file system on a DASD hard disk. The device must already have had a low-level format.

Format



-b *blocksize*

Specifies the blocksize. Default is 4096. The blocksize needs to be a multiple of the blocksize specified on the **dasdfmt** command. This allows you to use block sizes larger than the hardware maximum of 4096.

Examples

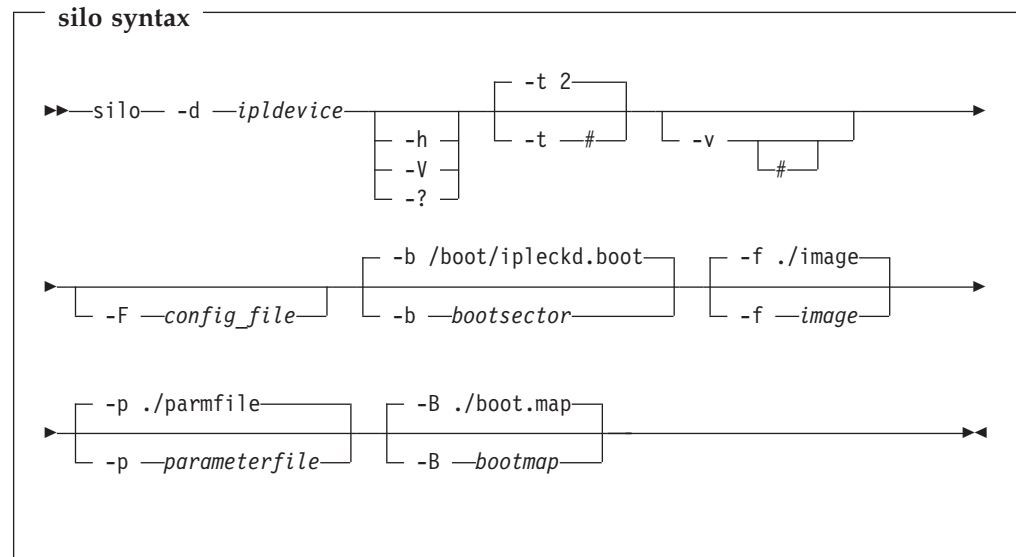
```
mke2fs -b 4096 /dev/dasda1
```

silo - Make DASD bootable

Usage

This tool is used to make DASDs (direct access storage devices) bootable. It takes a kernel image, a parameter file, a boot sector file, and the device node as input. Additionally, the file `/etc/silo.conf`, or the file specified by using the `-F` file name parameter, is parsed for additional options.

Format



Note that the defaults for these parameters can be overwritten by entering keywords in the config-file. The format used for each parameter keyword is shown in monospaced text in the following descriptions.

-d ipldevice

Set `ipldevice=devicenode` to set the boot device to a specific device node. The device node specified must be the node of the 'full' device and not one of a partition.

-h Prints out a short usage message.

-V Print version number and exit silo.

-? Prints out a short usage message. Equivalent to `-h`

-t[#] By default, silo runs with a `testlevel` of 2, which means that no modifications are made to the disk. A testing level of 1 means that a bootmap is generated with a temporary file name, but the boot records of the disk are not modified. The disk is made bootable only with a testing level of 0 or below. Set `testlevel=level` to decrease the testing level from the default by the value of *level*. Use the short forms `-t` to decrease the testing level by one, or `-t#`, to decrease the testing level by #.

-v[#] Set `verbose=level` to specify the level of verbosity used. Increases verbosity, or sets verbosity to #, respectively.

-F config-file

There are some defaults for the most common parameters compiled into

the binary. You can overwrite these defaults by your own using `/etc/silo.conf` or another config-file specified by `-F config-file`. All values set by defaults or the config-file can be overwritten using the command line options of `silo`.

An example of the contents of this file is:

```
image = image ipldevice = /dev/dasdb parmfile = parmfile bootsect = ipleckd.boot testlevel =
```

-b bootsector

Set `bootsect=bootsect` to specify the name of the bootsector to be used as boot record for that volume. The default name is `/boot/ipleckd.boot`.

-f image

Set `image=image` to specify the name of the image that is going to be loaded from that volume. The default name is `./image`.

-p parameterfile

Set `parmfile=parameter file` to specify the name of the parameter file holding the kernel parameters to be used during setup of the kernel. The default name is `./parmfile`.

-B bootmap

Set `map=bootmap` to specify the name of the bootmap used to hold the map information needed during booting. The default name is `./boot.map`. In test-only mode this name is replaced by a temporary name.

Chapter 10. Kernel parameters

There are different ways of passing parameters to LINUX :

- Passing parameters to your kernel at start-up time. (The parameter line)
- Configuring your boot loader to always pass those parameters.


The kernel can only handle a parmline file that is no larger than 896 bytes.

ipldelay

Usage

When you do a power on reset, some activation and loading is done. This can cause LINUX not to find the OSA-2 card. If you have problems with your OSA-2 card after booting, you might want to insert a delay to allow the POR, microcode load and initialization to take place in the OSA-2 card. The recommended delay time is two minutes. For example, 30s means a delay of thirty seconds between the boot and the initialization of the OSA-2 card, 2m means a delay of two minutes. The value xyz must be a number followed by either s or m.

Format



Examples

This example delays the initialization of the card by 2 minutes:

```
ipldelay=2m
```

This example delays the initialization of the card by 30 seconds:

```
ipldelay=30s
```

maxcpus

Usage

Specifies the maximum number of CPUs that LINUX can use.

Format

maxcpus syntax

►► `maxcpus` = `number` ◄◄

Examples


`maxcpus=2`

mem

Usage

Restricts memory usage to the size specified. This must be used to overcome initialisation problems on a P/390.

Format



Examples

```
mem=64M
```

Restricts the memory LINUX can use to 64MB.

```
mem=123456K
```

Restricts the memory LINUX can use to 123456KB.

noinitrd

Usage

The `noinitrd` statement is required when the kernel was compiled with initial RAM disk support enabled. This command bypasses using the initial ramdisk.

This can be useful if the kernel was used with a RAM disk for the initial start-up, but the RAM disk is not required when booted from a DASD

Format

noinitrd syntax

►► `noinitrd` ◄◄

ramdisk_size

Usage

Specifies the size of the ramdisk in kilobytes.

Format

ramdisk_size syntax

▶▶ramdisk_size = *size*◀◀

Examples

ramdisk_size=32000

ro

Usage

Mounts the root file system read-only.

Format

ro syntax

►► ro ◀◀

root

Usage

Tells LINUX what to use as the root when mounting the root file system.

Format

root syntax

►►—root— = —*mountpoint*—◄◄

Examples

This makes LINUX use /dev/devda1 when mounting the root file system:

```
root=/dev/dasda1
```

vmhalt

Usage

Specifies a command to be issued after a shutdown on VM.

Format

vmhalt syntax

►►—vmhalt— = —*command*—◄◄

Examples

This example specifies that an initial program load of CMS should follow a shutdown on VM:

```
vmhalt="i cms"
```

Chapter 11. Overview of the parameter line file

Parameters

Comments

The parameters allowed in the parameter line are described in “Chapter 10. Kernel parameters” on page 65 and/or together with the description of the device drivers.

Usage

The parameter line file contains data to be passed to the kernel for evaluation at start-up time. The location from which the kernel reads this file varies with the IPL method as shown below:

IPL method	Location of parameter line file
DASD	Installed into the boot sector using <code>sil0</code> (option <code>-p</code>)
Tape	Second file on tape
VM reader	Second file in reader
CD-ROM	Third entry in the <code>.ins</code> file (with load address 0x00010480)

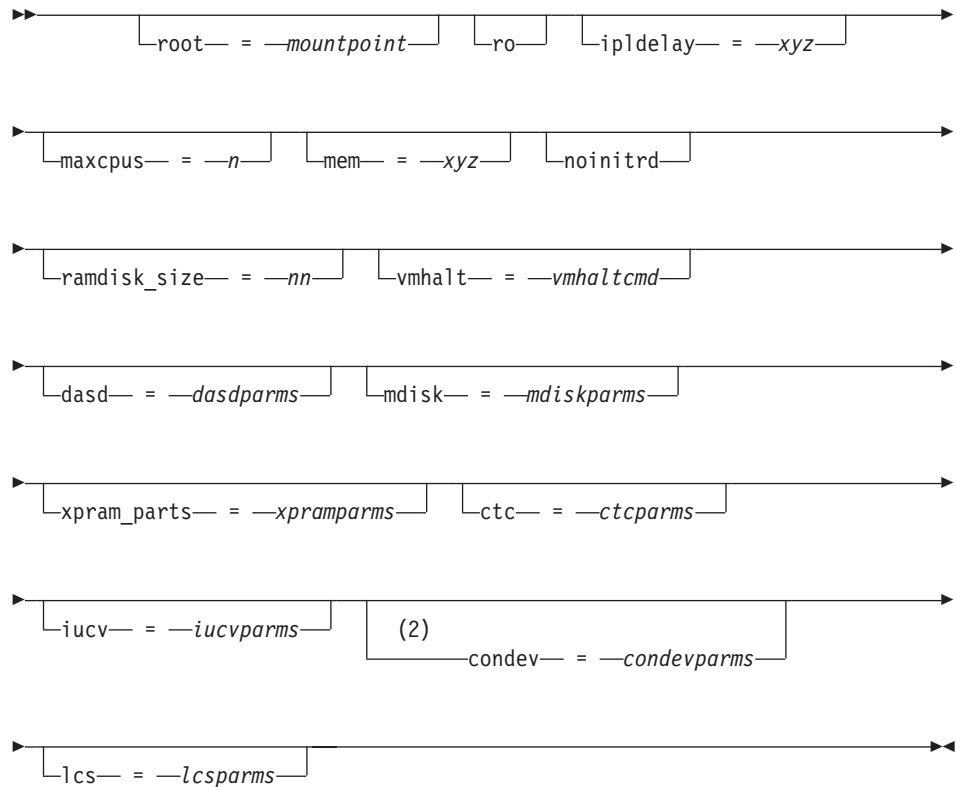
Format

The kernel parameter file consists of a single line containing at most 896 bytes. The line may either be encoded in ASCII or in EBCDIC. It contains a list of kernel options (see kernel parameters, device driver parameters) separated by blanks.

For IPL from a VM reader the kernel parameter file must be broken into fixed length records of 80 bytes. Note that a record end does not separate two options. Therefore if an option ends at the end of a record the next record should begin with a blank character.

Parameter line syntax

(1)



Notes:

- 1 The order of the parameters is irrelevant.
- 2 P/390 only

Where:

root=mountpoint

is defined in “root” on page 72

ro is defined in “ro” on page 71

ipldelay=xyz

is defined in “ipldelay” on page 66

dasdparms

are defined in “Chapter 1. LINUX for S/390 DASD device driver” on page 5

mdiskparms

are defined in “Chapter 2. LINUX for S/390 VM minidisk device driver” on page 13

xpramparms

are defined in “Chapter 3. LINUX for S/390 XPRAM device driver” on page 17

ctcparms

are defined in “Chapter 4. LINUX for S/390 CTC/ESCON device driver” on page 19

iucvparms

are defined in “Chapter 5. LINUX for S/390 IUCV device driver” on page 25

condevparms

are defined in “Chapter 6. LINUX for S/390 Console device drivers” on page 31

lcsparms

are defined in “Chapter 7. LINUX for S/390 LCS Device Driver” on page 37

Examples

Here is an example of the content of a parameter line file:

```
dasd=E0C0-E0C2 root=/dev/ram0 ro ipldelay=2m
```

Note that when loading from tape using an ASCII encoded parameter file (such as one generated on a UNIX or PC system) you must make sure that your parameter file does not span more than one line, is not larger than 896 bytes, and contains no special characters (for example tabs or new lines).

Part 5. Appendixes

Appendix A. Reference information

LCS module parameter syntax	81	Gigabit Ethernet driver command syntax	81
LCS kernel parameter syntax	81	LINUX for S/390 Device numbers.	82

LCS module parameter syntax

This driver is subject to licence conditions as reflected in: "International License Agreement for Non-Warranted Programs" on page 115.

The following are the LCS device driver module parameters:

use_hw_stats	Get network statistics from LANSTAT LCS primitive.
do_sw_ip_checksumming	Perform checksum on inbound packets.
additional_model_info	Model/maximum relative adapter number pairs.
devno_portno_pairs	Matching pairs of device numbers and port numbers.
noauto	noauto=1 disables auto-detection.

For a description of the parameters see "LCS module parameter syntax" on page 39.

LCS kernel parameter syntax

This driver is subject to licence conditions as reflected in: "International License Agreement for Non-Warranted Programs" on page 115.

If the LCS driver is compiled directly into the kernel the LCS boot parameters are as follows:

hw_stats	Get network statistics from LANSTAT LCS primitive.
ip_checksumming	Perform checksum on inbound packets.
additional cu model, max rel adapter no	Model/maximum relative adapter number pairs.
devno,rel_adapter_no	Matching pairs of device numbers and port numbers.
lcs_noauto	disables auto-detection.

For a description of the parameters see "LCS kernel parameter syntax" on page 38.

Gigabit Ethernet driver command syntax

This driver is subject to licence conditions as reflected in: "International License Agreement for Non-Warranted Programs" on page 115.

There is a single keyword parameter for the Gigabit Ethernet driver:

qeth_options

This parameter is used as follows:

(Note that all characters must be in the case shown, except in hexadecimal numbers where case is irrelevant.)

```
qeth_options=[<driver options>],[<card options>],[<card options>,...]]
```

<driver options> is a comma separated list that sets the driver defaults. It can contain the following keywords:

auto	turns on autosensing
noauto	turns off autosensing
no_router	does not prepare the device as router (default)
primary_router	makes the device a primary router
secondary_router	makes the device a secondary router
sw_checksumming	checksumming is to be performed by the software
hw_checksumming	checksumming is to be performed by the hardware
no_checksumming	checksumming is not to be used
prio_queueing_tos	priority queueing based on the IP type of service field
prio_queueing_prec	priority queueing based on the IP precedence field
no_prio_queueing	switch off priority queueing
no_prio_queueing: <number>	switch off priority queueing and set the default queue to <number>

<card options> are used to override the global options for a particular device. These are also comma-separated lists and each list consists of three device numbers (decimal, or hex starting with 0x), an optional device name, and any of the driver options keywords except **auto** or **noauto**.

For a description of the parameters see “GbE module parameter syntax” on page 43.

LINUX for S/390 Device numbers

Device numbers currently allocated to S/390 devices are:

Device	Major number	Minor numbers
DASD	94 + dynamic	0,4,8..252. – Volume, 1,5,9,...253 – Partitions
VM Minidisk	95	0–255
XPRAM	35	0–31

Appendix B. Kernel building

Building the kernel	83	Partition types	100
Preliminary steps	84	Kernel hacking	101
Configuring the parameters	85	Load an alternate configuration file	101
Checking the configuration	86	Save configuration to an alternate file	102
Checking the dependencies	86	Exit 'menuconfig'	102
Compiling the kernel	86	Kernel parameter options	103
Compiling for IPL from tape.	87	IEEE FPU emulation	103
Installing the modules.	87	Built-in IPL record support.	104
Finishing off	87	IPL method generated into head.S	104
Using 'config' or 'oldconfig'	88	Support for VM minidisk	104
Sample output listing	88	Support for VM minidisk synchronous	
Cross-reference to configuration options	90	read-write	104
Using 'menuconfig'	91	Support for DASD devices	105
File handling	91	Support for ECKD disks.	105
Main menu	93	Support for FBA disks	105
Code maturity level option	94	Support for DIAG access to CMS reserved	
Processor type and features	94	minidisk	106
Loadable module support	95	XPRAM device support	106
General setup	95	CTC/ESCON device support	106
S/390 block device drivers	96	IUCV device support.	107
S/390 network device support	97	Support for 3215 line mode terminal	107
S/390 terminal and console options	97	Support for console output on 3215 line mode	
Networking options	98	terminal	107
QoS and/or Fair queueing	99	Support for hardware console	108
Filesystems	99	Console output on hardware console	108
Network file systems	100		

Building the kernel

Before deciding to change the details in the kernel source code, consider whether installing one of the LINUX for S/390 kernel images provided in the LINUX for S/390 kernel patches will be a more appropriate solution to your requirements.

Your build system must have the following software installed (as a minimum):

- kernel source 2.2.16 with the LINUX for S/390 patch
- gcc version 2.95.2 or newer supporting LINUX for S/390
- binutils version 2.9.1 or newer supporting LINUX for S/390
- glibc 2.1.2 or newer supporting LINUX for S/390
- sed
- bash
- make version 3.77 or newer.

The following assumptions are made:

- You are confident in your ability to modify the kernel parameters without severely damaging the system

- You cannot locate a pre-compiled kernel image that meets your requirements (that is, a suitable kernel does not already exist)
- You are able to make an emergency IPL tape available (or preferably a complete backup on tape).

If you decide to modify your LINUX for S/390 kernel, you should use the instructions outlined in the following sections. In this way you will be sure of completing all of the tasks necessary to ensure the system runs properly when you have finished. For example, you might have to install and link additional modules after you have compiled and installed the kernel.

1. "Preliminary steps"
2. "Configuring the parameters" on page 85
3. "Checking the configuration" on page 86
4. "Checking the dependencies" on page 86
5. "Compiling the kernel" on page 86
6. "Installing the modules" on page 87
7. "Finishing off" on page 87

Preliminary steps

Before working with the kernel, there are a number of precautions that you should take:

- Make a backup copy of the current kernel and all modules corresponding to this kernel. It is important to make a backup even if you are replacing your kernel with a new version. This is because the new system might not run properly and you can use the backup to reload the old system
- Decide whether you want to modify the complete kernel, or only change some modules. If you only change some modules, you might not have to build the kernel.

If you are upgrading or replacing the kernel, obtain the new kernel or patch and load it into the directory `/usr/src`. This will probably create a new directory `usr/src/linux`, which will overwrite your last version of the kernel source. You will need to check the symbolic links to the `/usr/include` directory to ensure the following two links are valid:

- `ln -sf /usr/src/linux/include/linux /usr/include/linux`
- `ln -sf /usr/src/linux/include/asm /usr/include/asm`

Your first step in modifying the kernel, is to change to the `/usr/src/linux` directory and enter the command `make distclean`. This cleans up the LINUX for S/390 distribution, resetting all options to their default values and removing all object files from the system. If you enter `make clean` instead of `make distclean`, you will only delete the object files.

Configuring the parameters

To configure the parameters, make sure you are in the `/usr/src/linux` directory and enter the command `make config`.

In `make config`, you select what you want to include in the resident kernel and what features you want to have available as dynamically loadable modules. See “Using ‘config’ or ‘oldconfig’” on page 88 for an example of a `make config` listing. You will generally select the minimal resident set that is needed to boot:

- The type of file system used for your root partition (for example, `ext2`)
- Normal hard disk drive support (for example, `DASD`)
- Network support
- TCP/IP support.

The set of modules is constantly increasing, and you will be able to select the option (M) when responding to the prompts shown in `make config` for those features that the current kernel can offer as loadable modules. You can completely enable or disable the use of your set of modules by using the `CONFIG_MODVERSIONS` option during `make config`.

`make config` requires `bash` to allow it work. `Bash` will be searched for in `$BASH`, `/bin/bash` and `/bin/sh` (in that order), so it must be located in one of these directories for it to work. `make config` must be performed even if you are only upgrading to the next patch. New configuration options are added in each release, and odd problems will turn up if the configuration files are not set up as expected. If you want to upgrade your existing configuration with minimal work, use `make oldconfig`, which will keep your old kernel and only ask you questions about new or modified options.

Alternative configuration commands are:

- `make menuconfig` — Text based menus, radiolists and windows, see “Using ‘menuconfig’” on page 91
- `make oldconfig` — Same as `make config` except all questions based on the contents of your existing `./.config` file
- `make xconfig` — This X windows based configuration tool is currently not available with LINUX for S/390.

Notes on make config:

- Keeping unnecessary drivers in the LINUX for S/390 kernel will make it bigger, and can cause problems: for example, unnecessary networking options might confuse some drivers.
- Selecting the kernel hacking option and changing the source code directly usually result in a bigger or slower LINUX for S/390 kernel (or both). Thus you should probably answer (N) to the questions for development, experimental, or debugging features.

Checking the configuration

There are a pair of scripts that check the source tree for problems. These scripts do not have to be run each time you build the kernel, but it is a good idea to check for these types of errors and discrepancies at regular intervals.

make checkconfig checks the source tree for missing instances of `#include <linux>`. This needs to be done occasionally, because the C preprocessor will silently give bad results if these symbols haven't been included (it treats undefined symbols in preprocessor directives as defined to 0). Superfluous uses of `#include <linux>` are also reported, but you can ignore these, because smart `CONFIG_*` dependencies make them harmless. You can run make checkconfig without configuring the kernel. Also, make checkconfig does not modify any files.

make checkhelp checks the source tree for options that are in `Config.in` files but are not documented in `scripts/Configure.help`. Again, this needs to be done occasionally. If you have hacked the kernel and changed configuration options or are adding new ones, it is a good idea to make checkhelp (and add help as necessary) before you publish your patch. Also, make checkhelp does not modify any files.

Checking the dependencies

All of the source dependencies must be set each time you configure a new LINUX for S/390 kernel.

Enter make dep to set up all the dependencies correctly. make dep is a synonym for the long form, make depend. This command performs two tasks:

- It computes dependency information about which `.o` files depend on which `.h` files. It records this information in a top-level file named `.hdepend` and in one file per source directory named `.depend`.
- If you have `CONFIG_MODVERSIONS` enabled, make dep computes symbol version information for all of the files that export symbols (note that both resident and modular files can export symbols). If you do not enable `CONFIG_MODVERSIONS`, you only have to run make dep once, right after the first time you configure the kernel. The `.hdepend` files and the `.depend` file are independent of your configuration. If you do enable `CONFIG_MODVERSIONS`, you must run make dep because the symbol version information depends on the configuration.

Compiling the kernel

Enter make image to create a LINUX for S/390 kernel image. This compiles the source code and leaves the kernel image in the current directory (usually `/usr/src/linux/arch/s390/boot`).

Note that make zImage and make bzImage are not supported by the LINUX for S/390 kernel.

Compiling for IPL from tape

If you want to make a boot tape, you must transfer a set of files to the IPL tape. The files, `image.tape.bin` (renamed as `image.txt`), `parm.line`, and `initrd.bin` (renamed as `initrd.txt`) are the ones used during the installation process.

If you want to IPL from tape, ensure that the following configuration settings are used:

- `CONFIG_IPL` is set
- `CONFIG_IPL_TAPE` is set
- `CONFIG_BLK_DEV_RAM` is set
- `CONFIG_BLK_DEV_INITRD` is set.

Additionally you should keep the default configuration settings to make sure that all requirements to get a running LINUX for S/390 kernel are met.

Installing the modules

If you configured any of the parts of the LINUX for S/390 kernel as modules by selecting (M) in the kernel parameter option, you will have to create the modules and then link them to the kernel.

You create the modules by entering the command `make modules`. This will compile all of the modules and update the `linux/modules` directory. This directory will now contain a set of symbolic links, pointing to the various object files in the kernel tree.

After you have created all your modules, you must enter `make modules_install`. This will copy all of the newly made modules into subdirectories under `/lib/modules/kernel_release/`, where `kernel_release` is 2.2.16 (the current kernel version).

As soon as you have rebooted the newly made kernel, you can use the utilities `insmod` and `rmmod` to install and remove modules without recompiling the kernel. Read the man-pages for `insmod` and `rmmod` to find out how to configure and remove a module.

Finishing off

You should always keep a backup LINUX for S/390 kernel available in case something goes wrong. This backup must also include the modules corresponding to that kernel. If you are installing a new kernel with the same version number as your working kernel, make a backup of your modules' directory before you do a `make modules_install`.

In order to boot your new kernel, you'll need to copy the kernel image (found in `/usr/src/linux/arch/s390/boot/image` after compilation) to the place where your regular bootable kernel is located. This will be on your IPL tape.

To see how to create a tape from which you can perform an IPL, refer to *LINUX for S/390 Installation, Configuration and Use*

Using 'config' or 'oldconfig'

Use make config to configure all of the LINUX for S/390 kernel options, or use make oldconfig to keep your current kernel options and change only those options that are new or have been modified in the latest kernel release.

Use make config or make oldconfig when you only have access to a line based console. If you can use a screen based console, you might find make menuconfig gives you a better interface (see "Using 'menuconfig'" on page 91).

The following output listing shows an example of the complete configuration script that results from a make config. See "Kernel parameter options" on page 103 for more information about individual options.

Sample output listing

```
rm -f include/asm
( cd include ; ln -sf asm-s390 asm)
/bin/sh scripts/Configure arch/s390/config.in
#
# Using defaults found in .config
#
*
* Code maturity level options
*
Prompt for development and/or incomplete code/drivers (CONFIG_EXPERIMENTAL) [Y/n/?]
*
* Processor type and features
*
Symmetric multi-processing support (CONFIG_SMP) [Y/n/?]
IEEE FPU emulation (CONFIG_IEEEFPU_EMULATION) [Y/n/?]
*
* Loadable module support
*
Enable loadable module support (CONFIG_MODULES) [Y/n/?]
Set version information on all symbols for modules (CONFIG_MODVERSIONS) [N/y/?]
Kernel module loader (CONFIG_KMOD) [Y/n/?]
*
* General setup
*
Fast IRQ handling (CONFIG_FAST_IRQ) [Y/n/?]
Built-in IPL record support (CONFIG_IPL) [Y/n/?]
IPL method generated into head.S (tape, vm_reader) [vm_reader]
defined CONFIG_IPL_VM
Networking support (CONFIG_NET) [Y/n/?]
System V IPC (CONFIG_SYSVIPC) [Y/n/?]
BSD Process Accounting (CONFIG_BSD_PROCESS_ACCT) [N/y/?]
Sysctl support (CONFIG_SYSCTL) [Y/n/?]
Kernel support for ELF binaries (CONFIG_BINFMT_ELF) [Y/m/n/?]
*
* S/390 block device drivers
*
Loopback device support (CONFIG_BLK_DEV_LOOP) [Y/m/n/?]
Network block device support (CONFIG_BLK_DEV_NBD) [N/y/m/?]
Multiple devices driver support (CONFIG_BLK_DEV_MD) [N/y/?]
RAM disk support (CONFIG_BLK_DEV_RAM) [Y/m/n/?]
Initial RAM disk (initrd) support (CONFIG_BLK_DEV_INITRD) [Y/n/?]
XPRAM disk support (CONFIG_BLK_DEV_XPRAM) [N/y/m/?]
Support for VM minidisk (VM only) (CONFIG_MDISK) [N/y/?] y
Support for synchronous read-write (CONFIG_MDISK_SYNC) [N/y/?] (NEW)
Support for DASD devices (CONFIG_DASD) [Y/m/n/?]
```

```

*
* DASD disciplines
*
    Support for ECKD Disks (CONFIG_DASD_ECKD) [Y/n/?]
    Support for FBA Disks (CONFIG_DASD_FBA) [Y/n/?]
    Support for S/390 tape devices (CONFIG_S390_TAPE) [N/y/m/?]
    Support for DIAG access to CMS reserved minidisk (CONFIG_DASD_MDSK) [N/y/?]
*
* S/390 Network device support
*
    Channel Device Configuration (Temporary Option) (CONFIG_CHANDEV) [N/y/?]
    Network device support (CONFIG_NETDEVICES) [Y/n/?]
    scripts/Configure: menu_option: command not found
*
* S/390 Network devices
*
    LAN Channel Station Interface (CONFIG_LCS) [Y/m/n/?]
    CTC device support (CONFIG_CTC) [Y/n/?]
    IUCV device support (VM only) (CONFIG_IUCV) [Y/n/?]
    Dummy net driver support (CONFIG_DUMMY) [N/y/m/?]
    Ethernet (10 or 100Mbit) (CONFIG_NET_ETHERNET) [Y/n/?]
    Token Ring driver support (CONFIG_TR) [Y/n/?]
*
* S/390 Terminal and Console options
*
    Support for 3215 line mode terminal (CONFIG_3215) [Y/n/?]
    Support for console on 3215 line mode terminal (CONFIG_3215_CONSOLE) [Y/n/?]
    Support for HWC line mode terminal (CONFIG_HWC) [Y/n/?]
    console on HWC line mode terminal (CONFIG_HWC_CONSOLE) [Y/n/?]
*
* Networking options
*
    Packet socket (CONFIG_PACKET) [Y/m/n/?]
    Kernel/User netlink socket (CONFIG_NETLINK) [Y/n/?]
    Routing messages (CONFIG_RTNETLINK) [N/y/?]
    Netlink device emulation (CONFIG_NETLINK_DEV) [N/y/m/?]
    Network firewalls (CONFIG_FIREWALL) [N/y/?]
    Socket Filtering (CONFIG_FILTER) [N/y/?]
    Unix domain sockets (CONFIG_UNIX) [Y/m/n/?]
    TCP/IP networking (CONFIG_INET) [Y/n/?]
    IP: multicasting (CONFIG_IP_MULTICAST) [N/y/?]
    IP: advanced router (CONFIG_IP_ADVANCED_ROUTER) [N/y/?]
    IP: kernel level autoconfiguration (CONFIG_IP_PNP) [N/y/?]
    IP: optimize as router not host (CONFIG_IP_ROUTER) [N/y/?]
    IP: tunneling (CONFIG_NET_IPIP) [N/y/m/?]
    IP: GRE tunnels over IP (CONFIG_NET_IPGRE) [N/y/m/?]
    IP: aliasing support (CONFIG_IP_ALIAS) [N/y/?]
    IP: TCP syncookie support (not enabled per default) (CONFIG_SYN_COOKIES) [N/y/?]
*
* (it is safe to leave these untouched)
*
    IP: Reverse ARP (CONFIG_INET_RARP) [N/y/m/?]
    IP: Allow large windows (not recommended if <16Mb of memory) (CONFIG_SKB_LARGE) [Y/n/?]
    The IPv6 protocol (EXPERIMENTAL) (CONFIG_IPV6) [N/y/m/?]
*
*
*
    The IPX protocol (CONFIG_IPX) [N/y/m/?]
    Appletalk DDP (CONFIG_ATALK) [N/y/m/?]
    CCITT X.25 Packet Layer (EXPERIMENTAL) (CONFIG_X25) [N/y/m/?]
    LAPB Data Link Driver (EXPERIMENTAL) (CONFIG_LAPB) [N/y/m/?]
    Bridging (EXPERIMENTAL) (CONFIG_BRIDGE) [N/y/?]
    802.2 LLC (EXPERIMENTAL) (CONFIG_LLC) [N/y/?]
    Acorn Econet/AUN protocols (EXPERIMENTAL) (CONFIG_ECONET) [N/y/m/?]
    WAN router (CONFIG_WAN_ROUTER) [N/y/m/?]
    Fast switching (read help!) (CONFIG_NET_FASTROUTE) [N/y/?]
    Forwarding between high speed interfaces (CONFIG_NET_HW_FLOWCONTROL) [N/y/?]

```

```

CPU is too slow to handle full bandwidth (CONFIG_CPU_IS_SLOW) [N/y/?]
*
* QoS and/or fair queueing
*
QoS and/or fair queueing (CONFIG_NET_SCHED) [N/y/?]
*
* Filesystems
*
Quota support (CONFIG_QUOTA) [N/y/?]
Kernel automounter support (CONFIG_AUTOFS_FS) [N/y/m/?]
ADFS filesystem support (read only) (EXPERIMENTAL) (CONFIG_ADFS_FS) [N/y/m/?]
Amiga FFS filesystem support (CONFIG_AFFS_FS) [N/y/m/?]
Apple Macintosh filesystem support (experimental) (CONFIG_HFS_FS) [N/y/m/?]
DOS FAT fs support (CONFIG_FAT_FS) [N/y/m/?]
ISO 9660 CDROM filesystem support (CONFIG_ISO9660_FS) [N/y/m/?]
Minix fs support (CONFIG_MINIX_FS) [N/y/m/?]
NTFS filesystem support (read only) (CONFIG_NTFS_FS) [N/y/m/?]
OS/2 HPFS filesystem support (read only) (CONFIG_HPFS_FS) [N/y/m/?]
/proc filesystem support (CONFIG_PROC_FS) [Y/n/?]
QNX filesystem support (EXPERIMENTAL) (CONFIG_QNX4FS_FS) [N/y/m/?]
ROM filesystem support (CONFIG_ROMFS_FS) [N/y/m/?]
Second extended fs support (CONFIG_EXT2_FS) [Y/m/n/?]
System V and Coherent filesystem support (CONFIG_SYSV_FS) [N/y/m/?]
UFS filesystem support (CONFIG_UFS_FS) [N/y/m/?]
SGI EFS filesystem support (read only) (experimental) (CONFIG_EFS_FS) [N/y/m/?]
*
* Network File Systems
*
Coda filesystem support (advanced network fs) (CONFIG_CODA_FS) [N/y/m/?]
NFS filesystem support (CONFIG_NFS_FS) [Y/m/n/?]
NFS server support (CONFIG_NFSD) [N/y/m/?]
SMB filesystem support (to mount WfW shares etc.) (CONFIG_SMB_FS) [N/y/m/?]
NCP filesystem support (to mount NetWare volumes) (CONFIG_NCP_FS) [N/y/m/?]
*
* Partition Types
*
BSD disklabel (BSD partition tables) support (CONFIG_BSD_DISKLABEL) [N/y/?]
Macintosh partition map support (CONFIG_MAC_PARTITION) [N/y/?]
SMD disklabel (Sun partition tables) support (CONFIG_SMD_DISKLABEL) [N/y/?]
Solaris (x86) partition table support (CONFIG_SOLARIS_X86_PARTITION) [N/y/?]
Unixware slices support (EXPERIMENTAL) (CONFIG_UNIXWARE_DISKLABEL) [N/y/?]
*
* Kernel hacking
*
Kernel profiling support (CONFIG_PROFILE) [N/y/?]
Remote GDB kernel debugging (CONFIG_REMOTE_DEBUG) [N/y/?]

*** End of Linux kernel configuration.
*** Check the top-level Makefile for additional configuration.
*** Next, you may run 'make dep'.

```

Cross-reference to configuration options

The LINUX for S/390-specific configuration options shown in the sample output listing are described in the following sections:

- “IEEE FPU emulation” on page 103
- “Built-in IPL record support” on page 104
- “IPL method generated into head.S” on page 104
- “Support for VM minidisk” on page 104
- “Support for VM minidisk synchronous read-write” on page 104
- “Support for DASD devices” on page 105
- “Support for ECKD disks” on page 105
- “Support for FBA disks” on page 105

- “Support for DIAG access to CMS reserved minidisk” on page 106
- “XPRAM device support” on page 106
- “CTC/ESCON device support” on page 106
- “IUCV device support” on page 107
- “Support for 3215 line mode terminal” on page 107
- “Support for console output on 3215 line mode terminal” on page 107
- “Support for hardware console” on page 108
- “Console output on hardware console” on page 108

Using 'menuconfig'

You can use make menuconfig to edit individual LINUX for S/390 kernel options out of sequence and you can discard your settings at any time. You need a screen based console device to use make menuconfig. If you only have access to a line based console, you must use make config, see “Using 'config' or 'oldconfig'” on page 88 for more details.

Some options can be built directly into the kernel. Other options can be made into dynamically loadable modules. Some options can be completely removed altogether. There are also certain kernel parameters which are not really selectable options (Y) or (N), but instead values must be entered for them or selected from a list (decimal or hexadecimal numbers or possibly text).

The menu items begin with various types of bracket to indicate that the features:

- [*] are configured to be built in to the kernel.
- [] are configured to be removed from the kernel.
- < M > are configured to be modularized.
- < > are module capable features.
- < * > could have been modules, but have been built into the kernel.

Modules are linked to the kernel after booting.

Some menu items contain subordinate options that are displayed only when the menu item has been selected.

File handling

You can revise your settings up until you save the configuration. You will be given a last chance to confirm them prior to exiting menuconfig, see “Exit 'menuconfig'” on page 102.

If menuconfig quits with an error while saving your configuration, you can look in the file /usr/src/linux/.menuconfig.log for information which can help you determine the failure cause.

You can also save the current configuration to a file of your choice, or load a previously saved configuration, see “Save configuration to an alternate file” on page 102 and “Load an alternate configuration file” on page 101.

Menuconfig supports the use of alternate configuration files for those who, for various reasons, find it necessary to switch between different kernel configurations. At the end of the main menu you will find two options. One is for saving the current configuration to a file of your choosing. The other option is for loading a previously saved alternate configuration. Even if you don't use alternate

configuration files, but during a session you decide to restore your previously saved settings from `.config`, you can use the Load Alternate... to do so without restarting `menuconfig`.

Main menu

The following example shows the different sets of configuration options:

Linux Kernel v2.2.16 Configuration

```
----- Main Menu -----
Arrow keys navigate the menu.  <Enter> selects submenus --->.
Highlighted letters are hotkeys.  Pressing <Y> includes, <N> excludes,
<M> modularizes features.  Press <Esc><Esc> to exit, <?> for Help.
Legend: [*] built-in  [ ] excluded  <M> module  < > module capable

Code maturity level options --->
Processor type and features --->
Loadable module support --->
General setup --->
S/390 block device drivers --->
S/390 Network device support --->
S/390 Terminal and Console options --->
--- Character devices
[*] Unix98 PTY support
(256) Maximum number of Unix98 PTYs in use (0-2048)
Networking options --->
Filesystems --->
Kernel hacking --->
---
Load an Alternate Configuration File
Save Configuration to an Alternate File

<Select>  < Exit >  < Help >
```

The options on this menu are described in detail in the following sections:

- “Code maturity level option” on page 94
- “Processor type and features” on page 94
- “Loadable module support” on page 95
- “General setup” on page 95
- “S/390 block device drivers” on page 96
- “S/390 network device support” on page 97
- “S/390 terminal and console options” on page 97
- “Networking options” on page 98
- “QoS and/or Fair queueing” on page 99
- “Filesystems” on page 99
- “Network file systems” on page 100
- “Partition types” on page 100
- “Kernel hacking” on page 101
- “Save configuration to an alternate file” on page 102
- “Load an alternate configuration file” on page 101
- “Exit ‘menuconfig’” on page 102.

Code maturity level option

The Code Maturity Level option is intended to reduce the number of options that are displayed. This can help by not showing the code/drivers that are either incomplete or still under development:

Linux Kernel v2.2.16 Configuration

Code maturity level options

Arrow keys navigate the menu. <Enter> selects submenus --->. Highlighted letters are hotkeys. Pressing <Y> includes, <N> excludes, <M> modularizes features. Press <Esc><Esc> to exit, <?> for Help. Legend: [*] built-in [] excluded <M> module < > module capable

[*] Prompt for development and/or incomplete code/drivers

<Select> < Exit > < Help >

Processor type and features

The Processor Type and Features options allow you to configure the kernel to match your S/390 hardware installation.

Linux Kernel v2.2.16 Configuration

Processor type and features

Arrow keys navigate the menu. <Enter> selects submenus --->. Highlighted letters are hotkeys. Pressing <Y> includes, <N> excludes, <M> modularizes features. Press <Esc><Esc> to exit, <?> for Help. Legend: [*] built-in [] excluded <M> module < > module capable

[*] Symmetric multi-processing support
[*] IEEE FPU emulation

<Select> < Exit > < Help >

The LINUX for S/390 Processor Type and Features menu option is described in the following section:

- “IEEE FPU emulation” on page 103.

Loadable module support

The Loadable Module Support option lets you use dynamically loaded modules that are linked to your basic kernel.

Linux Kernel v2.2.16 Configuration

Loadable module support

Arrow keys navigate the menu. <Enter> selects submenus --->. Highlighted letters are hotkeys. Pressing <Y> includes, <N> excludes, <M> modularizes features. Press <Esc><Esc> to exit, <?> for Help. Legend: [*] built-in [] excluded <M> module < > module capable

[*] Enable loadable module support
[] Set version information on all symbols for modules
[*] Kernel module loader

<Select> < Exit > < Help >

General setup

The General Setup options configure your kernel's interaction with certain external programs, this includes:

- kernel networking support
- the synchronization and exchange of information between kernel and program
- instructing the kernel to write process accounting information to a file
- dynamically changing certain kernel parameters and variables on the fly.

Linux Kernel v2.2.16 Configuration

General setup

Arrow keys navigate the menu. <Enter> selects submenus --->. Highlighted letters are hotkeys. Pressing <Y> includes, <N> excludes, <M> modularizes features. Press <Esc><Esc> to exit, <?> for Help. Legend: [*] built-in [] excluded <M> module < > module capable

[*] Fast IRQ handling
[*] Built-in IPL record support
(tape) IPL method generated into head.S
[*] Networking support
[*] System V IPC
[] BSD Process Accounting
[] Sysctl support
<*> Kernel support for ELF binaries

<Select> < Exit > < Help >

The LINUX for S/390 General Setup menu options are described in the following sections:

- “Built-in IPL record support” on page 104
- “IPL method generated into head.S” on page 104.

S/390 block device drivers

The S/390 Block Device Drivers options allow the kernel to be set up to use the various block devices available with your S/390 system.

Linux Kernel v2.2.16 Configuration

S/390 block device_drivers

Arrow keys navigate the menu. <Enter> selects submenus --->.
Highlighted letters are hotkeys. Pressing <Y> includes, <N> excludes,
<M> modularizes features. Press <Esc><Esc> to exit, <?> for Help.
Legend: [*] built-in [] excluded <M> module < > module capable

```

<*> Loopback device support
< > Network block device support
[ ] Multiple devices driver support
<*> RAM disk support
[*]   Initial RAM disk (initrd) support
< > XPRAM disk support
[*] Support for VM minidisk (VM only)
[ ]   Support for synchronous read-write
<*> Support for DASD devices
--- DASD disciplines
[*]   Support for ECKD Disks
[*]   Support for FBA Disks
< > Support for S/390 tape devices
[ ] Support for DIAG access to CMS reserved minidisk

```

<Select> < Exit > < Help >

Note that the Support for DIAG access to CMS reserved minidisk option is available only when the Support for VM minidisk option is not selected.

The S/390 Block Device Drivers menu options unique to LINUX for S/390 are described in the following sections:

- “Support for VM minidisk” on page 104
- “Support for VM minidisk synchronous read-write” on page 104
- “Support for DASD devices” on page 105
- “Support for ECKD disks” on page 105.
- “Support for FBA disks” on page 105
- “Support for DIAG access to CMS reserved minidisk” on page 106
- “XPRAM device support” on page 106

S/390 network device support

The S/390 Network Device Support options allow the kernel to be set up to use the various network devices available with your S/390 system.

Linux Kernel v2.2.16 Configuration

```
----- S/390 Network device support -----
Arrow keys navigate the menu.  <Enter> selects submenus --->.
Highlighted letters are hotkeys.  Pressing <Y> includes, <N> excludes,
<M> modularizes features.  Press <Esc><Esc> to exit, <?> for Help.
Legend: [*] built-in  [ ] excluded  <M> module  < > module capable

[*] Network device support
--- S390 Network devices
<*> Lan Channel Station Interface
[*] CTC device support
[*] IUCV device support (VM only)
< > Dummy net driver support
[*] Ethernet (10 or 100Mbit)
[*] Token Ring driver support

<Select>    < Exit >    < Help >
```

The S/390 Network device support menu options unique to LINUX for S/390 are described in the following sections:

- “CTC/ESCON device support” on page 106
- “IUCV device support” on page 107.

S/390 terminal and console options

The S/390 Terminal and Console options allow the kernel to be set up to use the various line mode terminals (or emulators) available with your S/390 system.

Linux Kernel v2.2.16 Configuration

```
----- S/390 Terminal and Console options -----
Arrow keys navigate the menu.  <Enter> selects submenus --->.
Highlighted letters are hotkeys.  Pressing <Y> includes, <N> excludes,
<M> modularizes features.  Press <Esc><Esc> to exit, <?> for Help.
Legend: [*] built-in  [ ] excluded  <M> module  < > module capable

[*] Support for 3215 line mode terminal
[*] Support for console on 3215 line mode terminal
[*] Support for hardware console
[*] console output on hardware console

<Select>    < Exit >    < Help >
```

The S/390 Terminal and Console options menu options unique to LINUX for S/390 are described in the following sections:

- “Support for 3215 line mode terminal” on page 107
- “Support for console output on 3215 line mode terminal” on page 107
- “Support for hardware console” on page 108
- “Console output on hardware console” on page 108.

Networking options

The Networking options allow the kernel to interact with the network devices attached to your S/390 and also lets you optimize the performance of communications within and outside your system.

Linux Kernel v2.2.16 Configuration

Networking options

Arrow keys navigate the menu. <Enter> selects submenus --->.
Highlighted letters are hotkeys. Pressing <Y> includes, <N> excludes,
<M> modularizes features. Press <Esc><Esc> to exit, <?> for Help.
Legend: [*] built-in [] excluded <M> module < > module capable

```

<*> Packet socket
[*] Kernel/User netlink socket
[ ] Routing messages
< > Netlink device emulation
[ ] Network firewalls
[ ] Socket Filtering
<*> Unix domain sockets
[*] TCP/IP networking
[ ] IP: multicasting
[ ] IP: advanced router
[ ] IP: kernel level autoconfiguration
[ ] IP: optimize as router not host
< > IP: tunneling
< > IP: GRE tunnels over IP
[ ] IP: aliasing support
[ ] IP: TCP syncookie support (not enabled per default)
--- (it is safe to leave these untouched)
< > IP: Reverse ARP
[*] IP: Allow large windows (not recommended if <16Mb of memory)
< > The IPv6 protocol (EXPERIMENTAL)
---
< > The IPX protocol
< > Appletalk DDP
< > CCITT X.25 Packet Layer (EXPERIMENTAL)
< > LAPB Data Link Driver (EXPERIMENTAL)
[ ] Bridging (EXPERIMENTAL)
[ ] 802.2 LLC (EXPERIMENTAL)
< > Acorn Econet/AUN protocols (EXPERIMENTAL)
< > WAN router
[ ] Fast switching (read help!)
[ ] Forwarding between high speed interfaces
[ ] CPU is too slow to handle full bandwidth
QoS and/or fair queueing --->

```

<Select> < Exit > < Help >

QoS and/or Fair queueing

The QoS and/or Fair Queueing option allows fine tuning of the network device packet handling algorithms.

Linux Kernel v2.2.16 Configuration

QoS and/or fair queueing

Arrow keys navigate the menu. <Enter> selects submenus --->. Highlighted letters are hotkeys. Pressing <Y> includes, <N> excludes, <M> modularizes features. Press <Esc><Esc> to exit, <?> for Help. Legend: [*] built-in [] excluded <M> module < > module capable

[] QoS and/or fair queueing

<Select> < Exit > < Help >

Filesystems

The Filesystems options allow you to define the filesystems used to access your storage devices (hard disk, CDROM etc.).

Linux Kernel v2.2.16 Configuration

Filesystems

Arrow keys navigate the menu. <Enter> selects submenus --->. Highlighted letters are hotkeys. Pressing <Y> includes, <N> excludes, <M> modularizes features. Press <Esc><Esc> to exit, <?> for Help. Legend: [*] built-in [] excluded <M> module < > module capable

[] Quota support
< > Kernel automounter support
< > ADFS filesystem support (read only) (EXPERIMENTAL)
< > Amiga FFS filesystem support
< > Apple Macintosh filesystem support (experimental)
< > DOS FAT fs support
< > ISO 9660 CDROM filesystem support
< > Minix fs support
< > NTFS filesystem support (read only)
< > OS/2 HPFS filesystem support (read only)
[*] /proc filesystem support
< > QNX filesystem support (EXPERIMENTAL)
< > ROM filesystem support
<*> Second extended fs support
< > System V and Coherent filesystem support
< > UFS filesystem support
< > SGI EFS filesystem support (read only) (experimental)
Network File Systems --->
Partition Types --->

<Select> < Exit > < Help >

Network file systems

The Network File Systems options let you decide which network filesystem is most appropriate for your system.

Linux Kernel v2.2.16 Configuration

Network File Systems

Arrow keys navigate the menu. <Enter> selects submenus --->. Highlighted letters are hotkeys. Pressing <Y> includes, <N> excludes, <M> modularizes features. Press <Esc><Esc> to exit, <?> for Help. Legend: [*] built-in [] excluded <M> module < > module capable

< > Coda filesystem support (advanced network fs)
<*> NFS filesystem support
< > NFS server support
< > SMB filesystem support (to mount WfW shares etc.)
< > NCP filesystem support (to mount NetWare volumes)

<Select> < Exit > < Help >

Partition types

The Partition Types options lets you gain access to the hard disk partitions of other device types.

Linux Kernel v2.2.16 Configuration

Partition Types

Arrow keys navigate the menu. <Enter> selects submenus --->. Highlighted letters are hotkeys. Pressing <Y> includes, <N> excludes, <M> modularizes features. Press <Esc><Esc> to exit, <?> for Help. Legend: [*] built-in [] excluded <M> module < > module capable

[] BSD disklabel (BSD partition tables) support
[] Macintosh partition map support
[] SMD disklabel (Sun partition tables) support
[] Solaris (x86) partition table support
[] Unixware slices support (EXPERIMENTAL)

<Select> < Exit > < Help >

Kernel hacking

The Kernel Hacking options allow you access and control over the kernel in certain situations. These options should be used with care!

Linux Kernel v2.2.16 Configuration

Kernel hacking

Arrow keys navigate the menu. <Enter> selects submenus --->. Highlighted letters are hotkeys. Pressing <Y> includes, <N> excludes, <M> modularizes features. Press <Esc><Esc> to exit, <?> for Help. Legend: [*] built-in [] excluded <M> module < > module capable

[] Kernel profiling support
[] Remote GDB kernel debugging

<Select> < Exit > < Help >

Load an alternate configuration file

For various reasons, you might want to keep different kernel configurations available on a single S/390 system. Entering a file name here will allow you to later retrieve, modify and use the current configuration as an alternate to whatever configuration options you have selected at that time.

Enter the name of the configuration file you wish to load. Accept the name shown to restore the configuration you last retrieved. Leave blank to abort.

arch/s390/defconfig

< Ok > < Help >

Save configuration to an alternate file

For various reasons, you might want to keep several different kernel configurations available on a single S/390 system. If you have saved a previous configuration in a file other than the kernel's default, entering the name of the file here will allow you to modify that configuration.

Enter a filename to which this configuration should be saved as an alternate. Leave blank to abort.
<input type="text"/>
< Ok > < Help >

Exit 'menuconfig'

When you have completed your modifications to the kernel, you are asked whether you want to save the new kernel configuration. Normally, you would respond by selecting the Yes option, but you might have made modifications that you do not want to keep. In that case you can exit the configuration process without saving the changes by selecting the No option.

Do you wish to save your new kernel configuration?
< Yes > < No >

Kernel parameter options

The LINUX for S/390 specific kernel parameter options are described in the following sections:

- “IEEE FPU emulation”
- “Built-in IPL record support” on page 104
- “IPL method generated into head.S” on page 104
- “Support for VM minidisk” on page 104
- “Support for VM minidisk synchronous read-write” on page 104
- “Support for DASD devices” on page 105
- “Support for ECKD disks” on page 105
- “Support for FBA disks” on page 105
- “Support for DIAG access to CMS reserved minidisk” on page 106
- “XPRAM device support” on page 106
- “CTC/ESCON device support” on page 106
- “IUCV device support” on page 107
- “Support for 3215 line mode terminal” on page 107
- “Support for console output on 3215 line mode terminal” on page 107
- “Support for hardware console” on page 108
- “Console output on hardware console” on page 108.

IEEE FPU emulation

Configuration option

CONFIG_IEEEFPU_EMULATION

Capable of being a module? -- (Module Name)

No

Value Required by LINUX for S/390

Dependent on S/390 version, see Description

Description

With S/390 versions G5 and G6 (or newer), the S/390 systems are capable of calculating floating point numbers in IEEE-format.

LINUX for S/390 provides an IEEE floating point emulation. This can be configured on, enter (Y) or off, enter (N) during kernel compilation time. If you have IEEE-emulation configured on, floating point arithmetic can be performed on any S/390 system, however it will be slow on those systems using the emulation.

On S/390 versions G3, G4 (or elder ones) you must run with IEEE-emulation configured on (Y).

On S/390 versions G5, G6 (or newer) you can make use of the hardware IEEE-arithmetic. VM has to be enabled to allow it to use and provide the hardware IEEE-arithmetic. This occurs automatically when you are running VM 2.4 or you have VM 2.3 with a PTF applied that enables the IEEE-arithmetic. If you do not have VM 2.4 or did not apply the PTF to VM 2.3, then you still can (and have to) rely on the IEEE-emulation as on the older S/390 systems.

Built-in IPL record support

Configuration option
CONFIG_IPL

Capable of being a module? -- (Module Name)
No

Value Required by LINUX for S/390
Yes

Description
With this option turned on an IPL loader is generated at the start of the kernel image. That makes it possible to 'boot' from the kernel image directly without the need of a separate loader. This makes sense for a medium that is sequentially read from the start at IPL time like a (VM) reader or a tape. The type of the loader generated to the head of the kernel image is chosen by the 'IPL method generated into head.S' selection.

IPL method generated into head.S

Configuration option
CONFIG_IPL_VM

Capable of being a module? -- (Module Name)
No

Value Required by LINUX for S/390
See Description

Description
There are two loaders available for the generation into the kernel. 'tape' selects the loader for an IPL from a tape device, 'vm_reader' selects the loader for an IPL from a VM virtual reader.

Support for VM minidisk

Configuration option
CONFIG_MDISK

Capable of being a module? -- (Module Name)
No

Value Required by LINUX for S/390
No

Description
This option is used under VM only. With this flag enabled (Y) your S/390 can use a reserved minidisk under VM. VM internally uses the DIAG 250 to access the minidisk. When this boot parameter is enabled, several parameters (the virtual device ID, the size, offset and blocksize) must be passed to the kernel parameter file for each device. The minidisk must be formatted and reserved under VM/CMS.

When you are running a native installation, you would use CONFIG_DASD to configure your DASD.

Support for VM minidisk synchronous read-write

Configuration option
CONFIG_MDISK_SYNC

Capable of being a module? -- (Module Name)

No

Value Required by LINUX for S/390

No, VM only

Description

This option is used under VM only. You can make the minidisk operation synchronous. Normally a DIAG 250 is issued to start an I/O operation and the finish of the operation is reported with an external interrupt. With this flag enabled (Y), the DIAG 250 is issued synchronously.

Support for DASD devices

Configuration option

CONFIG_DASD

Capable of being a module? -- (Module Name)

dasd.o

Value Required by LINUX for S/390

See Description

Description

This is used mainly in native installations.

Enable this option (Y) to support access to S/390 disks. These are known as DASD (Direct Access Storage Device). You must enable this option to have disk access on a native or LPAR system. Running under VM you can choose CONFIG_MDISK instead.

When enabled (Y), this option lets you specify additional options for DASD access, see "Support for ECKD disks".

Support for ECKD disks

Configuration option

CONFIG_DASD_ECKD

Capable of being a module? -- (Module Name)

No

Value Required by LINUX for S/390

See Description

Description

This is used mainly in native installations.

Enable this option (Y) if you have ECKD-type DASDs like IBM 3380s or 3390s. ECKD devices are the most commonly used devices in S/390, so you should enable this option unless you are very sure you do not have any ECKD devices.

This option is a subordinate of CONFIG_DASD, see "Support for DASD devices".

Support for FBA disks

Configuration option

CONFIG_DASD_FBA

Capable of being a module? -- (Module Name)

No

Value Required by LINUX for S/390

See Description

Description

This is used mainly under VM/ESA for the virtual disk in storage VFB-512. Enable this option (Y) if you want to access your FBA devices.

This option is a sub-option of CONFIG_DASD, see "Support for DASD devices" on page 105.

Support for DIAG access to CMS reserved minidisk

Configuration option

CONFIG_DASD_MDSK

Capable of being a module? -- (Module Name)

No

Value Required by LINUX for S/390

See Description

Description

This is applicable under VM/ESA only. Enable this option (Y) if you want to access your disks by means of VM/ESA's DIAG250 opcode instead of channel processing. You might want to do this if channel access to your device is currently unsupported under LINUX for S/390 or you require the additional capabilities of VM/ESA such as cross-guest DASD caching or enhanced error recovery.

This option is a sub-option of CONFIG_DASD, see "Support for DASD devices" on page 105.

XPRAM device support

Configuration option

CONFIG_XPRAM

Capable of being a module? -- (Module Name)

xpram.o

Value Required by LINUX for S/390

See Description

Description

This is used to allow more than 2 GB of main storage to be accessed by LINUX for S/390. Enable this option (Y) to support access to expanded storage of up to 16 TB (S/390 hardware currently supports 64 GB). The expanded storage can be subdivided into partitions.

See "XPRAM kernel parameter syntax" on page 17 for more information.

CTC/ESCON device support

Configuration option

CONFIG_CTC

Capable of being a module? -- (Module Name)

lcs.o

Value Required by LINUX for S/390

No

Description

If you want to use channel connections under LINUX , enter (Y) here. This gives you the possibility to make TCP/IP connections via virtual, parallel or ESCON channels between LINUX for S/390 and other S/390 operating systems (LINUX for S/390, z/OS, OS/390, VM/ESA and VSE/ESA).

Read the Device Driver description for more information.

IUCV device support**Configuration option**

CONFIG_IUCV

Capable of being a module? -- (Module Name)

netiucv.o

Value Required by LINUX for S/390

No

Description

This is a VM only device driver. Enter (Y) to enable a fast communication between VM user IDs. At boot the VM user ID need to be passed to the kernel. Using ifconfig a point-to-point connection can be established to the LINUX for S/390 system running on the other VM user ID. Note that both kernels need to be compiled with this option and both need to be booted with the user ID of the other VM account.

Support for 3215 line mode terminal**Configuration option**

CONFIG_3215

Capable of being a module? -- (Module Name)

No

Value Required by LINUX for S/390

No

Description

The 3215 console driver is used to read and write to a 3215 line mode console. Real 3215 devices are no longer available in an S/390 environment, therefore the 3215 driver can only be used under VM. On a native S/390 system the initialization function of the 3215 driver returns without registering the driver to the system.

Entering (Y) to this option switches on the compilation of parts 1) and 2) of the 3215 terminal driver. The option makes it possible to use "Support for console output on 3215 line mode terminal".

Support for console output on 3215 line mode terminal**Configuration option**

CONFIG_3215_CONSOLE

Capable of being a module? -- (Module Name)

No

Value Required by LINUX for S/390

No

Description

This option is subordinate to "Support for 3215 line mode terminal".

This option enables console output on the first 3215 console in the system. It prints kernel errors and kernel warnings to the 3215 terminal in addition to the normal output on the TTY device.

Support for hardware console

Configuration option

CONFIG_HWC

Capable of being a module? -- (Module Name)

No

Value Required by LINUX for S/390

See Description

Description

The hardware console is an alternative terminal, usually required for a native LINUX for S/390 installation although it is also run under VM.

You would normally enter (Y) for this option in a native installation if your hardware configuration includes a hardware console. In a VM installation, without a hardware console, you would normally enter (N).

Read the Device Driver description for more information.

Console output on hardware console

Configuration option

CONFIG_HWC_CONSOLE

Capable of being a module? -- (Module Name)

No

Value Required by LINUX for S/390

No

Description

This option is subordinate to “Support for hardware console”.

This option enables console output on the first hardware console in the system. It prints kernel errors and kernel warnings to the hardware console in addition to the normal output on the TTY device.

Glossary

This glossary includes IBM product terminology as well as selected other terms and definitions.

Additional information can be obtained in:

- The American National Standard Dictionary for Information Systems , ANSI X3.172-1990, copyright 1990 by the American National Standards Institute (ANSI). Copies may be purchased from the American National Standards Institute, 11 West 42nd Street, New York, New York 10036.
- The ANSI/EIA Standard-440-A, Fiber Optic Terminology. Copies may be purchased from the Electronic Industries Association, 2001 Pennsylvania Avenue, N.W., Washington, DC 20006.
- The Information Technology Vocabulary developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC1/SC1).
- The IBM Dictionary of Computing , New York: McGraw-Hill, 1994.
- Internet Request for Comments: 1208, Glossary of Networking Terms
- Internet Request for Comments: 1392, Internet Users' Glossary
- The Object-Oriented Interface Design: IBM Common User Access Guidelines, Carmel, Indiana: Que, 1992.

A

autosensing. Listing the addresses of devices attached to a card by issuing a query command to the card.

C

checksum. An error detection method using a check byte appended to message data

CHPID. channel path identifier. In a channel subsystem, a value assigned to each installed channel path of the system that uniquely identifies that path to the system.

CRC. cyclic redundancy check. A system of error checking performed at both the sending and receiving station after a block-check character has been accumulated.

CTC. channel to channel. A method of connecting two computing devices.

CUU. Control Unit and Unit Address. A form of addressing for S/390 devices using device numbers.

D

DASD. direct access storage device. A mass storage medium on which a computer stores data.

device driver. (1) A file that contains the code needed to use an attached device. (2) A program that enables a computer to communicate with a specific peripheral device; for example, a printer, a videodisc player, or a CD-ROM drive. (3) A collection of subroutines that control the interface between I/O device adapters and the processor.

E

ECKD. extended count-key-data device. A disk storage device that has a data transfer rate faster than some processors can utilize and that is connected to the processor through use of a speed matching buffer. A specialized channel program is needed to communicate with such a device.

ESCON. Enterprise Systems Connection. A set of IBM products and services that provide a dynamically connected environment within an enterprise.

Ethernet. A 10-Mbps baseband local area network that allows multiple stations to access the transmission medium at will without prior coordination, avoids

contention by using carrier sense and deference, and resolves contention by using collision detection and delayed retransmission. Ethernet uses carrier sense multiple access with collision detection (CSMA/CD).

F

Fast Ethernet. Ethernet network with a bandwidth of 100-Mbps

FBA. fixed-block-architecture type DASD on P/390 (Multiprise 3000) or supported by VM.

FDDI. Fiber Distributed Data Interface. An American National Standards Institute (ANSI) standard for a 100-megabit-per-second LAN using optical fiber cables.

FTP. File Transfer Protocol. In the Internet suite of protocols, an application layer protocol that uses TCP and Telnet services to transfer bulk-data files between machines or hosts.

G

Gigabit Ethernet. Ethernet network with a bandwidth of 1000-Mbps

G3, G4, G5 and G6. The generation name of the S/390 CMOS based product family.

H

hardware console. This is a service-call logical processor that is the communication feature between the main processor and the service processor.

HMC. hardware management console. A console used to monitor and control hardware such as the S/390 microprocessors.

HFS. hierarchical file system.

I

IP. Internet Protocol. In the Internet suite of protocols, a connectionless protocol that routes data through a network or interconnected networks and acts as an intermediary between the higher protocol layers and the physical network.

IP address.. The unique 32-bit address that specifies the location of each device or workstation on the Internet. For example, 9.67.97.103 is an IP address.

IPL. initial program load (or boot). (1) The initialization procedure that causes an operating system to commence operation. (2) The process by which a configuration image is loaded into storage at the beginning of a work day or after a system malfunction.

(3) The process of loading system programs and preparing a system to run jobs.

IPv6. IP version 6. The next generation of the Internet Protocol.

IPX. Internetwork Packet Exchange. (1) The network protocol used to connect Novell's servers, or any workstation or router that implements IPX, with other workstations. Although similar to the Internet Protocol (IP), IPX uses different packet formats and terminology.

IPX address. IPX address. The 10-byte address, consisting of a 4-byte network number and a 6-byte node address, that is used to identify nodes in the IPX network. The node address is usually identical to the medium access control (MAC) address of the associated LAN adapter.

IUCV. inter-user communication vehicle. A VM facility for passing data between virtual machines and VM components.

K

kernel. The part of an operating system that performs basic functions such as allocating hardware resources.

kernel module. A dynamically loadable part of the kernel. It can be a device driver, file system, etc.

kernel image. The kernel when loaded into memory.

L

LCS. LAN channel station. A protocol used by OSA.

LDP. Linux Documentation Project. An attempt to provide a centralized location containing the source material for all open source LINUX documentation. Includes user and reference guides, HOW TOs, and FAQs.

LINUX for S/390. LINUX is a version of UNIX that runs on x86, Alpha and PowerPC machines. LINUX for S/390 is the port of LINUX to the IBM S/390 mainframe architecture.

LPAR. logical partition of an S/390

M

MAC. medium access control. In LANs, the sublayer of the data link control layer that supports medium-dependent functions and uses the services of the physical layer to provide services to the logical link control (LLC) sublayer. The MAC sublayer includes the method of determining when a device has access to the transmission medium.

MTU size. maximum transmission unit. The size of the largest block which may be transmitted as a single unit.

Multicast. A protocol for the simultaneous distribution of data to a number of recipients, for example live video transmissions.

Multiprise. An enterprise server of the S/390 family.

N

NIC. network interface card.

O

OS. operating system. (1) Software that controls the execution of programs. An operating system may provide services such as resource allocation, scheduling, input/output control, and data management. (2) A set of programs that control how the system works. (3) The software that deals with the most basic operations that a computer performs.

OSA-2. Open Systems Adapter-2. A common S/390 network interface card.

R

RAMAC.

router. A process which allows messages to pass between different networks.

S

S/390. System/390. Pertaining to the family of IBM enterprise servers that demonstrate outstanding reliability, availability, scalability, security, and capacity in today's network computing environments.

SA/SE. stand alone support element. See SE.

SE. support element. (1) An internal control element of a processor that assists in many of the processor operational functions. (2) A hardware unit that provides communications, monitoring, and diagnostic functions to a central processor complex.

SNA. Systems Network Architecture. The IBM architecture that defines the logical structure, formats, protocols, and operational sequences for transmitting information units through, and controlling the configuration and operation of, networks. The layered structure of SNA allows the ultimate origins and destinations of information (the users) to be independent of and unaffected by the specific SNA network services and facilities that are used for information exchange.

Sysctl. system control programming manual control (frame). A means of dynamically changing certain kernel parameters on the fly.

T

TCP. Transmission Control Protocol. A communications protocol used in the Internet and in any network that follows the Internet Engineering Task Force (IETF) standards for internetwork protocol. TCP provides a reliable host-to-host protocol between hosts in packet-switched communications networks and in interconnected systems of such networks. It uses the Internet Protocol (IP) as the underlying protocol.

TCP/IP. Transmission Control Protocol/Internet Protocol. (1) The Transmission Control Protocol and the Internet Protocol, which together provide reliable end-to-end connections between applications over interconnected networks of different types. (2) The suite of transport and application protocols that run over the Internet Protocol.

Telnet. In the Internet suite of protocols, a protocol that provides remote terminal connection service. It allows users of one host to log on to a remote host and interact as directly attached terminal users of that host.

Token Ring. (1) According to IEEE 802.5, network technology that controls media access by passing a token (special packet or frame) between media-attached stations. (2) A FDDI or IEEE 802.5 network with a ring topology that passes tokens from one attaching ring station (node) to another.

U

UNIX. UNIX operating system. An operating system developed by Bell Laboratories that features multiprogramming in a multiuser environment. The UNIX operating system was originally developed for use on minicomputers but has been adapted for mainframes and microcomputers.

V

volume. A data carrier that is mounted and demounted as a unit, for example, a reel of tape or a disk pack. Some disk units have no demountable packs. In that case, a volume is the portion available to one read/write mechanism.

Numbers

3215. IBM console printer-keyboard.

3270. IBM information display system.

3380 or 3390. IBM direct access storage device.

3480 or 3490. IBM magnetic tape subsystem.

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information about the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

Advanced Peer-to-Peer Networking	OSA
APPN	PowerPC
CICS	RACF
Common User Access	RAMAC
e-business	S/390
ECKD	Seascape
ESA/390	System/390
ESCON	VM/ESA
IBM	VSE/ESA
Micro Channel	VTAM
Multiprise	z/OS
MVS	zSeries
OS/2	
OS/390	

UNIX is a registered trademark in the United States and other countries licensed exclusively through The Open Group.

LINUX is a registered trademark of Linus Torvalds and others.

Microsoft, Windows NT, and MSDOS are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

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

International License Agreement for Non-Warranted Programs

- General Terms

The Program is owned by International Business Machines Corporation or one of its subsidiaries (IBM) or an IBM supplier, and is copyrighted and licensed, not sold.

The term "Program" means the original program and all whole or partial copies of it. A Program consists of machine-readable instructions, its components, data, audio-visual content (such as images, text, recordings, or pictures), and related licensed materials.

This Agreement is the complete agreement regarding the use of this Program, and replaces any prior oral or written communications between you and IBM.

1. Licence

- Use of the Program

IBM grants you a nonexclusive licence to use the Program.

You may 1) use the Program to the extent of authorizations you have acquired and 2) make, install and distribute copies to support the level of use authorized, providing you reproduce the copyright notice and any other legends of ownership on each copy, or partial copy, of the Program and that you provide a copy of this International License Agreement for Non-Warranted Programs along with the distribution of each copy or partial copy of the Program.

You will ensure that anyone who uses the Program does so only in compliance with the terms of this Agreement.

You may not 1) use, copy, modify, or distribute the Program except as provided in this Agreement; 2) reverse assemble, reverse compile, or otherwise translate the Program except as specifically permitted by law without the possibility of contractual waiver; or 3) sublicense, rent, or lease the Program.

2. No Warranty

SUBJECT TO ANY STATUTORY WARRANTIES WHICH CAN NOT BE EXCLUDED, IBM MAKES NO WARRANTIES OR CONDITIONS EITHER EXPRESS OR IMPLIED, INCLUDING WITHOUT LIMITATION, THE WARRANTY OF NON-INFRINGEMENT AND THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, REGARDING THE PROGRAM OR TECHNICAL SUPPORT, IF ANY. IBM MAKES NO WARRANTY REGARDING THE CAPABILITY OF THE PROGRAM TO CORRECTLY PROCESS, PROVIDE AND/OR RECEIVE DATE DATA WITHIN AND BETWEEN THE 20TH AND 21ST CENTURIES.

The exclusion also applies to any of IBM's subcontractors, suppliers, or program developers (collectively called "Suppliers").

Manufacturers, suppliers, or publishers of non-IBM Programs may provide their own warranties.

3. Limitation of Liability

NEITHER IBM NOR ITS SUPPLIERS WILL BE LIABLE FOR ANY DIRECT OR INDIRECT DAMAGES, INCLUDING WITHOUT LIMITATION, LOST PROFITS, LOST SAVINGS, OR ANY INCIDENTAL, SPECIAL, OR OTHER ECONOMIC CONSEQUENTIAL DAMAGES, EVEN IF IBM IS INFORMED OF THEIR POSSIBILITY. SOME JURISDICTIONS DO NOT ALLOW THE

EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THE ABOVE EXCLUSION OR LIMITATION MAY NOT APPLY TO YOU.

4. General

Nothing in this Agreement affects any statutory rights of consumers that cannot be waived or limited by contract.

IBM may terminate your licence if you fail to comply with the terms of this Agreement. If IBM does so, you must immediately destroy the Program and all copies you made of it.

You agree to comply with applicable export laws and regulations.

Neither you nor IBM will bring a legal action under this Agreement more than two years after the cause of action arose unless otherwise provided by local law without the possibility of contractual waiver or limitation.

Neither you nor IBM is responsible for failure to fulfill any obligations due to causes beyond its control.

IBM does not provide program services or technical support, unless IBM specifies otherwise.

GNU General Public Licence, Version 2, June 1991

DISCLAIMER

Elements of LINUX for S/390 which utilise internal details of the S/390 systems remain the intellectual property and copyright of IBM, notwithstanding the licence below. This right will be waived for specific elements in the case that the documentation specific to the element indicates that it is published under the GNU General Public Licence.

LINUX for S/390 is licensed under the GNU General Public Licence which is reproduced below:

Copyright (C) 1989, 1991
Free Software Foundation, Inc.
59 Temple Place,
Suite 330,
Boston,
MA 02111-1307
USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps:

1. copyright the software, and
2. offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU General Public Licence: Terms and conditions for copying, distribution and modification

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:
 - a. You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
 - b. You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

- c. If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

- 3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:
 - a. Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
 - b. Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
 - c. Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.
6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.
7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is

permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

<one line to give the program's name and a brief idea of what it does.>
Copyright (C) <year> <name of author>

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.
This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.
You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

Gnomovision version 69, Copyright (C) year name of author
Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type 'show w'.
This is free software, and you are welcome to redistribute it under certain conditions; type 'show c' for details.

The hypothetical commands 'show w' and 'show c' should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than 'show w' and 'show c'; they could even be mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the program, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the program 'Gnomovision' (which makes passes at compilers) written by James Hacker.
<signature of Ty Coon>, 1 April 1989
Ty Coon, President of Vice

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Library General Public License instead of this License.

Index

Numerics

3215 line mode terminal 31
3270 31
3380 7
3390 7
9345 7

A

autosensing 47, 82

B

basic mode 40

C

checksum 37, 39, 45, 82
chpid 47, 48
codepage 33
control characters 32
CRC 37, 39
CTC 37
 device driver 19
CTC recovery 24
CUU 37

D

DASD device driver 5
dasdfmt 54
device driver 43
 CTC 19
 DASD 5
 ESCON 19
 VM minidisk 13
 XPRAM 17
device name 45, 82
device number 45, 82

E

ECKD 7
Enterprise Storage Server 7
ESCON
 device driver 19
ethernet 37, 40

F

FBA 7
FDDI 40

G

Gigabit ethernet 43

H

Hardware console 31

I

ipldelay 66
IUCV 25

K

kernel 37, 38, 81
kernel source tree vii

L

LCS
 device driver 37
 driver parameters 39, 81

M

mac 41
maxcpus 67
mem 68
Multiprise 7

N

NFS 40
noinitrd 69
notices 113

O

OSA-2 37, 40
OSA-Express 43

P

P/390 31, 68
parameter line 76
problems 40

Q

qdio 43
QDIO 47, 50
qeth 43
qeth_options 81
queueing 43, 82

R

RAMAC 7
ro 71
root 72
routing 43, 45, 82
RVA 7

S

Seascape 7
silo 63

SMP 50
subchannel 47

T

TCP/IP 19, 25
token ring 37, 40
trademarks 114

V

VM minidisk
 device driver 13
vmhalt 73

X

x3270 33
XPRAM
 device driver 17



Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

LINUX-1003-02

