Full credit is given to the above companies including the OS that this PDF file was generated!

Red Hat Enterprise Linux Release 9.2 Manual Pages on 'sk98lin.4' command

$ man sk98lin.4

SK98LIN(4) Linux Programmer's Manual SK98LIN(4)

NAME

sk98lin - Marvell/SysKonnect Gigabit Ethernet driver v6.21

SYNOPSIS

insmod sk98lin.o [Speed_A=i,j,...] [Speed_B=i,j,...] [Au?
toNeg_A=i,j,...] [AutoNeg_B=i,j,...] [DupCap_A=i,j,...] [Dup?
Cap_B=i,j,...] [FlowCtrl_A=i,j,...] [FlowCtrl_B=i,j,...]
[Role_A=i,j,...] [Role_B=i,j,...] [ConType=i,j,...] [Modera?
tion=i,j,...] [IntsPerSec=i,j,...] [PrefPort=i,j,...] [Rlmt?
Mode=i,j,...]

DESCRIPTION

Note: This obsolete driver was removed from the kernel in version
2.6.26.

sk98lin is the Gigabit Ethernet driver for Marvell and SysKonnect
network adapter cards. It supports SysKonnect SK-98xx/SK-95xx
compliant Gigabit Ethernet Adapter and any Yukon compliant chipset.
When loading the driver using insmod, parameters for the network
adapter cards might be stated as a sequence of comma separated
commands. If for instance two network adapters are installed and
AutoNegotiation on Port A of the first adapter should be ON, but on the
Port A of the second adapter switched OFF, one must enter:

insmod sk98lin.o AutoNeg_A=On,Off

After sk98lin is bound to one or more adapter cards and the /proc

filesystem is mounted on your system, a dedicated statistics file will be created in the folder /proc/net/sk98lin for all ports of the installed network adapter cards. Those files are named eth[x], where x is the number of the interface that has been assigned to a dedicated port by the system.

If loading is finished, any desired IP address can be assigned to the respective eth[x] interface using the ifconfig(8) command. This causes the adapter to connect to the Ethernet and to display a status message on the console saying "ethx: network connection up using port y" followed by the configured or detected connection parameters.

The sk98lin also supports large frames (also called jumbo frames). Using jumbo frames can improve throughput tremendously when transferring large amounts of data. To enable large frames, the MTU (maximum transfer unit) size for an interface is to be set to a high value. The default MTU size is 1500 and can be changed up to 9000 (bytes). Setting the MTU size can be done when assigning the IP address to the interface or later by using the ifconfig(8) command with the mtu parameter. If for instance eth0 needs an IP address and a large frame MTU size, the following two commands might be used:

```
ifconfig eth0 10.1.1.1
ifconfig eth0 mtu 9000
```

Those two commands might even be combined into one:

```
ifconfig eth0 10.1.1.1 mtu 9000
```

Note that large frames can be used only if permitted by your network infrastructure. This means, that any switch being used in your Ethernet must also support large frames. Quite some switches support large frames, but need to be configured to do so. Most of the times, their default setting is to support only standard frames with an MTU size of 1500 (bytes). In addition to the switches inside the network, all network adapters that are to be used must also be enabled regarding jumbo frames. If an adapter is not set to receive large frames, it will simply drop them.

Switching back to the standard Ethernet frame size can be done by using

the ifconfig(8) command again:

ifconfig eth0 mtu 1500

The Marvell/SysKonnect Gigabit Ethernet driver for Linux is able to support VLAN and Link Aggregation according to IEEE standards 802.1, 802.1q, and 802.3ad. Those features are available only after installation of open source modules which can be found on the Internet:

VLAN: ?http://www.candelatech.com/~greear/vlan.html?

Link Aggregation: ?http://www.st.rim.or.jp/~yumo?

Note that Marvell/SysKonnect does not offer any support for these open source modules and does not take the responsibility for any kind of failures or problems arising when using these modules.

Parameters

Speed_A=i,j,...

This parameter is used to set the speed capabilities of port A of an adapter card. It is valid only for Yukon copper adapters. Possible values are: 10, 100, 1000, or Auto; Auto is the de? fault. Usually, the speed is negotiated between the two ports during link establishment. If this fails, a port can be forced to a specific setting with this parameter.

Speed_B=i,j,...

This parameter is used to set the speed capabilities of port B of an adapter card. It is valid only for Yukon copper adapters. Possible values are: 10, 100, 1000, or Auto; Auto is the de? fault. Usually, the speed is negotiated between the two ports during link establishment. If this fails, a port can be forced to a specific setting with this parameter.

AutoNeg_A=i,j,...

Enables or disables the use of autonegotiation of port A of an adapter card. Possible values are: On, Off, or Sense; On is the default. The Sense mode automatically detects whether the link partner supports auto-negotiation or not.

AutoNeg_B=i,j,...

Enables or disables the use of autonegotiation of port B of an

adapter card. Possible values are: On, Off, or Sense; On is the default. The Sense mode automatically detects whether the link partner supports auto-negotiation or not.

DupCap_A=i,j,...

This parameter indicates the duplex mode to be used for port A of an adapter card. Possible values are: Half, Full, or Both; Both is the default. This parameter is relevant only if Au? toNeg_A of port A is not set to Sense. If AutoNeg_A is set to On, all three values of DupCap_A (Half, Full or Both) might be stated. If AutoNeg_A is set to Off, only DupCap_A values Full and Half are allowed. This DupCap_A parameter is useful if your link partner does not support all possible duplex combinations.

DupCap_B=i,j,...

This parameter indicates the duplex mode to be used for port B of an adapter card. Possible values are: Half, Full, or Both; Both is the default. This parameter is relevant only if Au? toNeg_B of port B is not set to Sense. If AutoNeg_B is set to On, all three values of DupCap_B (Half, Full or Both) might be stated. If AutoNeg_B is set to Off, only DupCap_B values Full and Half are allowed. This DupCap_B parameter is useful if your link partner does not support all possible duplex combinations.

FlowCtrl_A=i,j,...

This parameter can be used to set the flow control capabilities the port reports during auto-negotiation. Possible values are: Sym, SymOrRem, LocSend, or None; SymOrRem is the default. The different modes have the following meaning:

Sym = Symmetric

both link partners are allowed to send PAUSE frames

SymOrRem = SymmetricOrRemote

both or only remote partner are allowed to send PAUSE frames

LocSend = LocalSend

only local link partner is allowed to send PAUSE frames

None = None

no link partner is allowed to send PAUSE frames

Note that this parameter is ignored if AutoNeg_A is set to Off.

FlowCtrl_B=i,j,...

This parameter can be used to set the flow control capabilities the port reports during auto-negotiation. Possible values are: Sym, SymOrRem, LocSend, or None; SymOrRem is the default. The different modes have the following meaning:

Sym = Symmetric

both link partners are allowed to send PAUSE frames

SymOrRem = SymmetricOrRemote

both or only remote partner are allowed to send PAUSE frames

LocSend = LocalSend

only local link partner is allowed to send PAUSE frames

None = None

no link partner is allowed to send PAUSE frames

Note that this parameter is ignored if AutoNeg_B is set to Off.

Role_A=i,j,...

This parameter is valid only for 1000Base-T adapter cards. For two 1000Base-T ports to communicate, one must take the role of the master (providing timing information), while the other must be the slave. Possible values are: Auto, Master, or Slave; Auto is the default. Usually, the role of a port is negotiated be? tween two ports during link establishment, but if that fails the port A of an adapter card can be forced to a specific setting with this parameter.

Role_B=i,j,...

This parameter is valid only for 1000Base-T adapter cards. For two 1000Base-T ports to communicate, one must take the role of the master (providing timing information), while the other must be the slave. Possible values are: Auto, Master, or Slave; Auto is the default. Usually, the role of a port is negotiated be? tween two ports during link establishment, but if that fails the port B of an adapter card can be forced to a specific setting

with this parameter.

ConType=i,j,...

This parameter is a combination of all five per-port parameters within one single parameter. This simplifies the configuration of both ports of an adapter card. The different values of this variable reflect the most meaningful combinations of port param? eters. Possible values and their corresponding combination of per-port parameters:

ConType	DupCap	AutoNeg	FlowCtrl	Role	Speed
Auto	Both	On	SymOrRem	Auto	Auto
100FD	Full	Off	None	Auto	100
100HD	Half	Off	None	Auto	100
10FD	Full	Off	None	Auto	10
10HD	Half	Off	None	Auto	10

Stating any other port parameter together with this ConType pa? rameter will result in a merged configuration of those settings. This is due to the fact, that the per-port parameters (e.g., Speed_A) have a higher priority than the combined variable Con? Type.

Moderation=i,j,...

Interrupt moderation is employed to limit the maximum number of interrupts the driver has to serve. That is, one or more inter? rupts (which indicate any transmit or receive packet to be pro? cessed) are queued until the driver processes them. When queued interrupts are to be served, is determined by the IntsPerSec pa? rameter, which is explained later below. Possible moderation modes are: None, Static, or Dynamic; None is the default. The different modes have the following meaning:

None No interrupt moderation is applied on the adapter card. Therefore, each transmit or receive interrupt is served immedi? ately as soon as it appears on the interrupt line of the adapter card.

Static Interrupt moderation is applied on the adapter card. All

transmit and receive interrupts are queued until a complete mod?
eration interval ends. If such a moderation interval ends, all
queued interrupts are processed in one big bunch without any de?
lay. The term Static reflects the fact, that interrupt modera?
tion is always enabled, regardless how much network load is cur?
rently passing via a particular interface. In addition, the du?
ration of the moderation interval has a fixed length that never
changes while the driver is operational.
Dynamic Interrupt moderation might be applied on the adapter
card, depending on the load of the system. If the driver de?
tects that the system load is too high, the driver tries to
shield the system against too much network load by enabling in?
terrupt moderation. If?at a later time?the CPU utilization de?
creases again (or if the network load is negligible), the inter?
rupt moderation will automatically be disabled.
Interrupt moderation should be used when the driver has to han?
dle one or more interfaces with a high network load, which?as a
consequence?leads also to a high CPU utilization. When modera?
tion is applied in such high network load situations, CPU load
might be reduced by 20?30% on slow computers.
Note that the drawback of using interrupt moderation is an in?
crease of the round-trip-time (RTT), due to the queuing and
serving of interrupts at dedicated moderation times.

IntsPerSec=i,j,...

This parameter determines the length of any interrupt moderation
interval. Assuming that static interrupt moderation is to be
used, an IntsPerSec parameter value of 2000 will lead to an in?
terrupt moderation interval of 500 microseconds. Possible val?
ues for this parameter are in the range of 30...40000 (inter?
rupts per second). The default value is 2000.
This parameter is used only if either static or dynamic inter?
rupt moderation is enabled on a network adapter card. This pa?
rameter is ignored if no moderation is applied.

Note that the duration of the moderation interval is to be cho?
sen with care. At first glance, selecting a very long duration
(e.g., only 100 interrupts per second) seems to be meaningful,
but the increase of packet-processing delay is tremendous. On
the other hand, selecting a very short moderation time might
compensate the use of any moderation being applied.

PrefPort=i,j,...

This parameter is used to force the preferred port to A or B (on
dual-port network adapters). The preferred port is the one that
is used if both ports A and B are detected as fully functional.
Possible values are: A or B; A is the default.

RlmtMode=i,j,...

RLMT monitors the status of the port. If the link of the active
port fails, RLMT switches immediately to the standby link. The
virtual link is maintained as long as at least one "physical"
link is up. This parameters states how RLMT should monitor both
ports. Possible values are: CheckLinkState, CheckLocalPort,
CheckSeg, or DualNet; CheckLinkState is the default. The dif?
ferent modes have the following meaning:
CheckLinkState Check link state only: RLMT uses the link state
reported by the adapter hardware for each individual port to de?
termine whether a port can be used for all network traffic or
not.
CheckLocalPort In this mode, RLMT monitors the network path be?
tween the two ports of an adapter by regularly exchanging pack?
ets between them. This mode requires a network configuration in
which the two ports are able to "see" each other (i.e., there
must not be any router between the ports).
CheckSeg Check local port and segmentation: This mode supports
the same functions as the CheckLocalPort mode and additionally
checks network segmentation between the ports. Therefore, this
mode is to be used only if Gigabit Ethernet switches are in?
stalled on the network that have been configured to use the

Spanning Tree protocol.

DualNet In this mode, ports A and B are used as separate de?
vices. If you have a dual port adapter, port A will be config?
ured as eth[x] and port B as eth[x+1]. Both ports can be used
independently with distinct IP addresses. The preferred port
setting is not used. RLMT is turned off.

Note that RLMT modes CheckLocalPort and CheckLinkState are de?
signed to operate in configurations where a network path between
the ports on one adapter exists. Moreover, they are not de?
signed to work where adapters are connected back-to-back.

FILES

/proc/net/sk98lin/eth[x]

The statistics file of a particular interface of an adapter
card. It contains generic information about the adapter card
plus a detailed summary of all transmit and receive counters.

/usr/src/linux/Documentation/networking/sk98lin.txt

This is the README file of the sk98lin driver. It contains a
detailed installation HOWTO and describes all parameters of the
driver. It denotes also common problems and provides the solu?
tion to them.

BUGS

Report any bugs to linux@syskonnect.de

SEE ALSO

ifconfig(8), insmod(8), modprobe(8)

COLOPHON

This page is part of release 5.10 of the Linux man-pages project. A
description of the project, information about reporting bugs, and the
latest version of this page, can be found at
https://www.kernel.org/doc/man-pages/.