

# **SCSI Interfaces Guide**

**James Bottomley**

**James.Bottomley@hansenpartnership.com**

**Rob Landley**

**rob@landley.net**

## **SCSI Interfaces Guide**

by James Bottomley and Rob Landley

Copyright © 2007 Linux Foundation

This documentation is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License version 2.

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**. For more details see the file **COPYING** in the source distribution of Linux.

# Table of Contents

<b>1. Introduction.....</b>	<b>1</b>
1.1. Protocol vs bus .....	1
1.2. Design of the Linux SCSI subsystem .....	1
<b>2. SCSI upper layer.....</b>	<b>3</b>
2.1. sd (SCSI Disk) .....	3
2.2. sr (SCSI CD-ROM).....	3
2.3. st (SCSI Tape) .....	3
2.4. sg (SCSI Generic) .....	3
2.5. ch (SCSI Media Changer).....	3
<b>3. SCSI mid layer .....</b>	<b>5</b>
3.1. SCSI midlayer implementation.....	5
3.1.1. include/scsi/scsi_device.h .....	5
shost_for_each_device .....	5
__shost_for_each_device .....	6
3.1.2. drivers/scsi/scsi.c.....	7
scsi_device_type .....	7
__scsi_get_command.....	7
scsi_get_command.....	8
__scsi_put_command.....	9
scsi_put_command.....	10
scsi_allocate_command .....	11
scsi_free_command.....	12
scsi_cmd_get_serial .....	13
scsi_finish_command.....	14
scsi_adjust_queue_depth.....	14
scsi_track_queue_full .....	16
scsi_get_vpd_page .....	17
scsi_device_get .....	18
scsi_device_put .....	19
starget_for_each_device.....	20
__starget_for_each_device.....	21
__scsi_device_lookup_by_target .....	22
scsi_device_lookup_by_target .....	23
__scsi_device_lookup .....	24
scsi_device_lookup .....	25
3.1.3. drivers/scsi/scsicam.c .....	26
scsi_bios_ptable .....	26
scsicam_bios_param .....	27
scsi_partsize .....	28
3.1.4. drivers/scsi/scsi_error.c .....	30

scsi_schedule_eh.....	30
scsi_block_when_processing_errors.....	31
scsi_eh_prep_cmnd.....	32
scsi_eh_restore_cmnd.....	33
scsi_eh_finish_cmd.....	34
scsi_eh_get_sense.....	35
scsi_eh_ready_devs.....	36
scsi_eh_flush_done_q.....	37
scsi_normalize_sense.....	37
scsi_sense_desc_find.....	39
scsi_get_sense_info fld.....	40
scsi_build_sense_buffer.....	41
3.1.5. drivers/scsi/scsi_devinfo.c.....	42
scsi_dev_info_list_add.....	42
scsi_dev_info_list_add_str.....	43
scsi_get_device_flags.....	44
scsi_exit_devinfo.....	45
scsi_init_devinfo.....	46
3.1.6. drivers/scsi/scsi_ioctl.c.....	46
scsi_ioctl.....	47
scsi_nonblockable_ioctl.....	48
3.1.7. drivers/scsi/scsi_lib.c.....	48
scsi_execute.....	49
scsi_device_from_queue.....	50
scsi_mode_select.....	51
scsi_mode_sense.....	53
scsi_test_unit_ready.....	54
scsi_device_set_state.....	55
sdev_evt_send.....	56
sdev_evt_alloc.....	57
sdev_evt_send_simple.....	58
scsi_device_quiesce.....	59
scsi_device_resume.....	60
scsi_internal_device_block.....	60
scsi_internal_device_unblock.....	61
scsi_kmap_atomic_sg.....	63
scsi_kunmap_atomic_sg.....	64
3.1.8. drivers/scsi/scsi_lib_dma.c.....	64
scsi_dma_map.....	64
scsi_dma_unmap.....	65
3.1.9. drivers/scsi/scsi_module.c.....	66
3.1.10. drivers/scsi/scsi_proc.c.....	66
proc_scsi_read.....	66

proc_scsi_write_proc .....	67
scsi_proc_hostdir_add.....	68
scsi_proc_hostdir_rm.....	69
scsi_proc_host_add.....	70
scsi_proc_host_rm .....	70
proc_print_scsidevice .....	71
scsi_add_single_device.....	72
scsi_remove_single_device.....	73
proc_scsi_write .....	74
proc_scsi_open.....	75
scsi_init_procfs .....	76
scsi_exit_procfs.....	77
3.1.11. drivers/scsi/scsi_netlink.c .....	78
scsi_nl_rcv_msg.....	78
scsi_nl_rcv_event.....	79
scsi_generic_msg_handler .....	79
scsi_netlink_init .....	80
scsi_netlink_exit .....	81
3.1.12. drivers/scsi/scsi_scan.c .....	82
scsi_complete_async_scans .....	82
scsi_unlock_floptical.....	83
scsi_alloc_sdev .....	84
scsi_alloc_target.....	85
scsi_target_reap.....	86
sanitize_inquiry_string.....	87
scsi_probe_lun .....	88
scsi_test_lun.....	89
scsi_add_lun.....	90
scsi_inq_str .....	91
scsi_probe_and_add_lun.....	92
scsi_sequential_lun_scan .....	94
scsi_report_lun_scan.....	95
scsi_prep_async_scan .....	96
scsi_finish_async_scan .....	97
3.1.13. drivers/scsi/scsi_sysctl.c .....	98
3.1.14. drivers/scsi/scsi_sysfs.c.....	98
scsi_remove_device .....	98
scsi_remove_target.....	99
3.1.15. drivers/scsi/hosts.c .....	99
scsi_host_set_state .....	99
scsi_remove_host.....	100
scsi_add_host_with_dma .....	101
scsi_host_alloc .....	102

scsi_host_lookup.....	103
scsi_host_get.....	104
scsi_host_put.....	104
scsi_queue_work.....	105
scsi_flush_work.....	106
3.1.16. drivers/scsi/constants.c.....	107
scsi_print_status.....	107
3.2. Transport classes .....	108
3.2.1. Fibre Channel transport.....	108
fc_get_event_number.....	108
fc_host_post_event.....	109
fc_host_post_vendor_event.....	110
fc_remove_host.....	111
fc_remote_port_add .....	112
fc_remote_port_delete .....	114
fc_remote_port_rolechg.....	115
fc_block_scsi_eh.....	117
fc_vport_create .....	118
fc_vport_terminate.....	118
3.2.2. iSCSI transport class .....	119
iscsi_create_flashnode_sess .....	119
iscsi_create_flashnode_conn.....	121
iscsi_is_flashnode_conn_dev .....	122
iscsi_find_flashnode_sess .....	123
iscsi_find_flashnode_conn .....	124
iscsi_destroy_flashnode_sess.....	125
iscsi_destroy_all_flashnode .....	126
iscsi_scan_finished.....	127
iscsi_block_scsi_eh.....	128
iscsi_unblock_session.....	128
iscsi_create_session .....	129
iscsi_destroy_session .....	130
iscsi_create_conn .....	131
iscsi_destroy_conn.....	132
iscsi_session_event .....	133
3.2.3. Serial Attached SCSI (SAS) transport class .....	134
sas_remove_children.....	134
sas_remove_host .....	135
sas_tlr_supported .....	136
sas_disable_tlr.....	136
sas_enable_tlr.....	137
sas_phy_alloc.....	138
sas_phy_add.....	139

sas_phy_free .....	140
sas_phy_delete .....	141
scsi_is_sas_phy .....	141
sas_port_add .....	142
sas_port_free .....	143
sas_port_delete.....	144
scsi_is_sas_port.....	145
sas_port_get_phy .....	146
sas_port_add_phy .....	146
sas_port_delete_phy.....	147
sas_end_device_alloc.....	148
sas_expander_alloc .....	149
sas_rphy_add.....	150
sas_rphy_free .....	151
sas_rphy_delete.....	152
sas_rphy_unlink .....	153
sas_rphy_remove .....	154
scsi_is_sas_rphy.....	154
sas_attach_transport.....	155
sas_release_transport .....	156
3.2.4. SATA transport class.....	157
3.2.5. Parallel SCSI (SPI) transport class .....	157
spi_schedule_dv_device .....	157
spi_display_xfer_agreement .....	158
3.2.6. SCSI RDMA (SRP) transport class .....	159
srp_rport_add .....	159
srp_rport_del .....	160
srp_remove_host .....	160
srp_attach_transport.....	161
srp_release_transport .....	162
<b>4. SCSI lower layer .....</b>	<b>165</b>
4.1. Host Bus Adapter transport types .....	165
4.1.1. Debug transport.....	165
4.1.2. todo .....	165



# Chapter 1. Introduction

## 1.1. Protocol vs bus

Once upon a time, the Small Computer Systems Interface defined both a parallel I/O bus and a data protocol to connect a wide variety of peripherals (disk drives, tape drives, modems, printers, scanners, optical drives, test equipment, and medical devices) to a host computer.

Although the old parallel (fast/wide/ultra) SCSI bus has largely fallen out of use, the SCSI command set is more widely used than ever to communicate with devices over a number of different busses.

The SCSI protocol (<http://www.t10.org/scsi-3.htm>) is a big-endian peer-to-peer packet based protocol. SCSI commands are 6, 10, 12, or 16 bytes long, often followed by an associated data payload.

SCSI commands can be transported over just about any kind of bus, and are the default protocol for storage devices attached to USB, SATA, SAS, Fibre Channel, FireWire, and ATAPI devices. SCSI packets are also commonly exchanged over Infiniband, I20 (<http://i2o.shadowconnect.com/faq.php>), TCP/IP (iSCSI (<http://en.wikipedia.org/wiki/ISCSI>)), even Parallel ports (<http://cyberelk.net/tim/parport/parscsi.html>).

## 1.2. Design of the Linux SCSI subsystem

The SCSI subsystem uses a three layer design, with upper, mid, and low layers. Every operation involving the SCSI subsystem (such as reading a sector from a disk) uses one driver at each of the 3 levels: one upper layer driver, one lower layer driver, and the SCSI midlayer.

The SCSI upper layer provides the interface between userspace and the kernel, in the form of block and char device nodes for I/O and `ioctl()`. The SCSI lower layer contains drivers for specific hardware devices.

In between is the SCSI mid-layer, analogous to a network routing layer such as the IPv4 stack. The SCSI mid-layer routes a packet based data protocol between the upper layer's `/dev` nodes and the corresponding devices in the lower layer. It manages command queues, provides error handling and power management functions, and responds to `ioctl()` requests.



# Chapter 2. SCSI upper layer

The upper layer supports the user-kernel interface by providing device nodes.

## 2.1. sd (SCSI Disk)

sd (sd\_mod.o)

## 2.2. sr (SCSI CD-ROM)

sr (sr\_mod.o)

## 2.3. st (SCSI Tape)

st (st.o)

## 2.4. sg (SCSI Generic)

sg (sg.o)

## 2.5. ch (SCSI Media Changer)

ch (ch.c)



# Chapter 3. SCSI mid layer

## 3.1. SCSI midlayer implementation

### 3.1.1. include/scsi/scsi\_device.h

## shost\_for\_each\_device

### LINUX

Kernel Hackers Manual July 2026

### Name

`shost_for_each_device` — iterate over all devices of a host

### Synopsis

```
shost_for_each_device ( sdev, shost );
```

### Arguments

*sdev*

the struct `scsi_device` to use as a cursor

*shost*

the struct `scsi_host` to iterate over

## Description

Iterator that returns each device attached to *shost*. This loop takes a reference on each device and releases it at the end. If you break out of the loop, you must call `scsi_device_put(sdev)`.

# \_\_shost\_for\_each\_device

## LINUX

Kernel Hackers Manual July 2026

## Name

`__shost_for_each_device` — iterate over all devices of a host (UNLOCKED)

## Synopsis

```
__shost_for_each_device ( sdev, shost );
```

## Arguments

*sdev*

the struct `scsi_device` to use as a cursor

*shost*

the struct `scsi_host` to iterate over

## Description

Iterator that returns each device attached to *shost*. It does `_not_` take a reference on the `scsi_device`, so the whole loop must be protected by `shost->host_lock`.

## Note

The only reason to use this is because you need to access the device list in interrupt context. Otherwise you really want to use `shost_for_each_device` instead.

### 3.1.2. drivers/scsi/scsi.c

Main file for the SCSI midlayer.

## scsi\_device\_type

### LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_device_type` — Return 17 char string indicating device type.

### Synopsis

```
const char * scsi_device_type (unsigned type);
```

### Arguments

*type*

type number to look up

# \_\_scsi\_get\_command

## LINUX

Kernel Hackers Manual July 2026

### Name

`__scsi_get_command` — Allocate a struct `scsi_cmnd`

### Synopsis

```
struct scsi_cmnd * __scsi_get_command (struct Scsi_Host *  
shost, gfp_t gfp_mask);
```

### Arguments

*shost*

host to transmit command

*gfp\_mask*

allocation mask

### Description

allocate a struct `scsi_cmnd` from host's slab, recycling from the host's `free_list` if necessary.

# scsi\_get\_command

## LINUX

## Name

`scsi_get_command` — Allocate and setup a scsi command block

## Synopsis

```
struct scsi_cmnd * scsi_get_command (struct scsi_device * dev,
gfp_t gfp_mask);
```

## Arguments

*dev*

parent scsi device

*gfp\_mask*

allocator flags

## Returns

The allocated scsi command structure.

# \_\_scsi\_put\_command

## LINUX

## Name

`__scsi_put_command` — Free a struct `scsi_cmnd`

## Synopsis

```
void __scsi_put_command (struct Scsi_Host * shost, struct  
scsi_cmnd * cmd, struct device * dev);
```

## Arguments

*shost*

dev->host

*cmd*

Command to free

*dev*

parent scsi device

## scsi\_put\_command

### LINUX

Kernel Hackers Manual July 2026

## Name

scsi\_put\_command — Free a scsi command block

## Synopsis

```
void scsi_put_command (struct scsi_cmnd * cmd);
```

## Arguments

*cmd*

command block to free

## Returns

Nothing.

## Notes

The command must not belong to any lists.

# scsi\_allocate\_command

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_allocate_command` — get a fully allocated SCSI command

## Synopsis

```
struct scsi_cmnd * scsi_allocate_command (gfp_t gfp_mask);
```

## Arguments

*gfp\_mask*

allocation mask

## Description

This function is for use outside of the normal host based pools. It allocates the relevant command and takes an additional reference on the pool it used. This function *must* be paired with `scsi_free_command` which also has the identical mask, otherwise the free pool counts will eventually go wrong and you'll trigger a bug.

This function should *only* be used by drivers that need a static command allocation at start of day for internal functions.

## scsi\_free\_command

### LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_free_command` — free a command allocated by `scsi_allocate_command`

### Synopsis

```
void scsi_free_command (gfp_t gfp_mask, struct scsi_cmnd *  
cmd);
```

### Arguments

*gfp\_mask*

mask used in the original allocation

*cmd*

command to free

## Note

using the original allocation mask is vital because that's what determines which command pool we use to free the command. Any mismatch will cause the system to BUG eventually.

# scsi\_cmd\_get\_serial

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_cmd_get_serial` — Assign a serial number to a command

## Synopsis

```
void scsi_cmd_get_serial (struct Scsi_Host * host, struct
scsi_cmnd * cmd);
```

## Arguments

*host*

the scsi host

*cmd*

command to assign serial number to

## Description

a serial number identifies a request for error recovery and debugging purposes. Protected by the `Host_Lock` of host.

## scsi\_finish\_command

### LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_finish_command` — cleanup and pass command back to upper layer

### Synopsis

```
void scsi_finish_command (struct scsi_cmnd * cmd);
```

### Arguments

*cmd*

the command

### Description

Pass command off to upper layer for finishing of I/O request, waking processes that are waiting on results, etc.

## scsi\_adjust\_queue\_depth

### LINUX

## Name

`scsi_adjust_queue_depth` — Let low level drivers change a device's queue depth

## Synopsis

```
void scsi_adjust_queue_depth (struct scsi_device * sdev, int tagged, int tags);
```

## Arguments

*sdev*

SCSI Device in question

*tagged*

Do we use tagged queueing (non-0) or do we treat this device as an untagged device (0)

*tags*

Number of tags allowed if tagged queueing enabled, or number of commands the low level driver can queue up in non-tagged mode (as per `cmd_per_lun`).

## Returns

Nothing

## Lock Status

None held on entry

## Notes

Low level drivers may call this at any time and we will do the right thing depending on whether or not the device is currently active and whether or not it even has the command blocks built yet.

# scsi\_track\_queue\_full

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_track_queue_full` — track QUEUE\_FULL events to adjust queue depth

## Synopsis

```
int scsi_track_queue_full (struct scsi_device * sdev, int depth);
```

## Arguments

*sdev*

SCSI Device in question

*depth*

Current number of outstanding SCSI commands on this device, not counting the one returned as QUEUE\_FULL.

## Description

This function will track successive QUEUE\_FULL events on a specific SCSI device to determine if and when there is a need to adjust the queue depth on the device.

## Returns

0 - No change needed, >0 - Adjust queue depth to this new depth, -1 - Drop back to untagged operation using host->cmd\_per\_lun as the untagged command depth

## Lock Status

None held on entry

## Notes

Low level drivers may call this at any time and we will do “The Right Thing.” We are interrupt context safe.

# scsi\_get\_vpd\_page

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_get_vpd_page` — Get Vital Product Data from a SCSI device

## Synopsis

```
int scsi_get_vpd_page (struct scsi_device * sdev, u8 page,  
unsigned char * buf, int buf_len);
```

## Arguments

*sdev*

The device to ask

*page*

Which Vital Product Data to return

*buf*

where to store the VPD

*buf\_len*

number of bytes in the VPD buffer area

## Description

SCSI devices may optionally supply Vital Product Data. Each 'page' of VPD is defined in the appropriate SCSI document (eg SPC, SBC). If the device supports this VPD page, this routine returns a pointer to a buffer containing the data from that page. The caller is responsible for calling `kfree` on this pointer when it is no longer needed. If we cannot retrieve the VPD page this routine returns `NULL`.

## scsi\_device\_get

### LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_device_get` — get an additional reference to a `scsi_device`

### Synopsis

```
int scsi_device_get (struct scsi_device * sdev);
```

## Arguments

*sdev*

device to get a reference to

## Description

Gets a reference to the `scsi_device` and increments the use count of the underlying LLDD module. You must hold `host_lock` of the parent `Scsi_Host` or already have a reference when calling this.

# scsi\_device\_put

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_device_put` — release a reference to a `scsi_device`

## Synopsis

```
void scsi_device_put (struct scsi_device * sdev);
```

## Arguments

*sdev*

device to release a reference on.

## Description

Release a reference to the `scsi_device` and decrements the use count of the underlying LLDD module. The device is freed once the last user vanishes.

# target\_for\_each\_device

## LINUX

Kernel Hackers Manual July 2026

## Name

`target_for_each_device` — helper to walk all devices of a target

## Synopsis

```
void target_for_each_device (struct scsi_target * target,  
void * data, void (*fn) (struct scsi_device *, void *));
```

## Arguments

*target*

target whose devices we want to iterate over.

*data*

Opaque passed to each function call.

*fn*

Function to call on each device

## Description

This traverses over each device of *starget*. The devices have a reference that must be released by `scsi_host_put` when breaking out of the loop.

# \_\_starget\_for\_each\_device

## LINUX

Kernel Hackers Manual July 2026

## Name

`__starget_for_each_device` — helper to walk all devices of a target (UNLOCKED)

## Synopsis

```
void __starget_for_each_device (struct scsi_target * starget,
void * data, void (*fn) (struct scsi_device *, void *));
```

## Arguments

*starget*

target whose devices we want to iterate over.

*data*

parameter for callback *fn*()

*fn*

callback function that is invoked for each device

## Description

This traverses over each device of *starget*. It does *\_not\_* take a reference on the *scsi\_device*, so the whole loop must be protected by *shost->host\_lock*.

## Note

The only reason why drivers would want to use this is because they need to access the device list in irq context. Otherwise you really want to use *target\_for\_each\_device* instead.

# \_\_scsi\_device\_lookup\_by\_target

## LINUX

Kernel Hackers Manual July 2026

## Name

`__scsi_device_lookup_by_target` — find a device given the target (UNLOCKED)

## Synopsis

```
struct scsi_device * __scsi_device_lookup_by_target (struct scsi_target * starget, uint lun);
```

## Arguments

*starget*

SCSI target pointer

*lun*

SCSI Logical Unit Number

## Description

Looks up the `scsi_device` with the specified `lun` for a given `target`. The returned `scsi_device` does not have an additional reference. You must hold the host's `host_lock` over this call and any access to the returned `scsi_device`. A `scsi_device` in state `SDEV_DEL` is skipped.

## Note

The only reason why drivers should use this is because they need to access the device list in irq context. Otherwise you really want to use `scsi_device_lookup_by_target` instead.

# scsi\_device\_lookup\_by\_target

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_device_lookup_by_target` — find a device given the target

## Synopsis

```
struct scsi_device * scsi_device_lookup_by_target (struct
scsi_target * target, uint lun);
```

## Arguments

*target*

SCSI target pointer

*lun*

SCSI Logical Unit Number

## Description

Looks up the `scsi_device` with the specified `lun` for a given `target`. The returned `scsi_device` has an additional reference that needs to be released with `scsi_device_put` once you're done with it.

# \_\_scsi\_device\_lookup

## LINUX

Kernel Hackers Manual July 2026

## Name

`__scsi_device_lookup` — find a device given the host (UNLOCKED)

## Synopsis

```
struct scsi_device * __scsi_device_lookup (struct Scsi_Host *  
shost, uint channel, uint id, uint lun);
```

## Arguments

*shost*

SCSI host pointer

*channel*

SCSI channel (zero if only one channel)

*id*

SCSI target number (physical unit number)

*lun*

SCSI Logical Unit Number

## Description

Looks up the `scsi_device` with the specified `channel`, `id`, `lun` for a given host. The returned `scsi_device` does not have an additional reference. You must hold the host's `host_lock` over this call and any access to the returned `scsi_device`.

## Note

The only reason why drivers would want to use this is because they need to access the device list in irq context. Otherwise you really want to use `scsi_device_lookup` instead.

# scsi\_device\_lookup

**LINUX**

Kernel Hackers Manual July 2026

## Name

`scsi_device_lookup` — find a device given the host

## Synopsis

```
struct scsi_device * scsi_device_lookup (struct Scsi_Host *  
shost, uint channel, uint id, uint lun);
```

## Arguments

*shost*

SCSI host pointer

*channel*

SCSI channel (zero if only one channel)

*id*

SCSI target number (physical unit number)

*lun*

SCSI Logical Unit Number

## Description

Looks up the `scsi_device` with the specified `channel`, `id`, `lun` for a given host. The returned `scsi_device` has an additional reference that needs to be released with `scsi_device_put` once you're done with it.

### 3.1.3. drivers/scsi/scsicam.c

SCSI Common Access Method

(<http://www.t10.org/ftp/t10/drafts/cam/cam-r12b.pdf>) support functions, for use with `HDIO_GETGEO`, etc.

# scsi\_bios\_ptable

## LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_bios_ptable` — Read PC partition table out of first sector of device.

### Synopsis

```
unsigned char * scsi_bios_ptable (struct block_device * dev);
```

### Arguments

*dev*

from this device

### Description

Reads the first sector from the device and returns 0x42 bytes starting at offset 0x1be.

### Returns

partition table in `kmalloc(GFP_KERNEL)` memory, or NULL on error.

# scsicam\_bios\_param

## LINUX

## Name

`scsicam_bios_param` — Determine geometry of a disk in cylinders/heads/sectors.

## Synopsis

```
int scsicam_bios_param (struct block_device * bdev, sector_t  
capacity, int * ip);
```

## Arguments

*bdev*

which device

*capacity*

size of the disk in sectors

*ip*

return value: `ip[0]`=heads, `ip[1]`=sectors, `ip[2]`=cylinders

## Description

determine the BIOS mapping/geometry used for a drive in a SCSI-CAM system, storing the results in `ip` as required by the `HDIO_GETGEO` `ioctl`.

## Returns

-1 on failure, 0 on success.

# scsi\_partsize

## LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_partsize` — Parse cylinders/heads/sectors from PC partition table

### Synopsis

```
int scsi_partsize (unsigned char * buf, unsigned long  
capacity, unsigned int * cyls, unsigned int * hds, unsigned  
int * secs);
```

### Arguments

*buf*

partition table, see `scsi_bios_ptable`

*capacity*

size of the disk in sectors

*cyls*

put cylinders here

*hds*

put heads here

*secs*

put sectors here

## Description

determine the BIOS mapping/geometry used to create the partition table, storing the results in `*cyls`, `*hds`, and `*secs`

## Returns

-1 on failure, 0 on success.

### 3.1.4. drivers/scsi/scsi\_error.c

Common SCSI error/timeout handling routines.

## scsi\_schedule\_eh

### LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_schedule_eh` — schedule EH for SCSI host

## Synopsis

```
void scsi_schedule_eh (struct Scsi_Host * shost);
```

## Arguments

*shost*

SCSI host to invoke error handling on.

## Description

Schedule SCSI EH without scmd.

# scsi\_block\_when\_processing\_errors

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_block_when_processing_errors` — Prevent cmds from being queued.

## Synopsis

```
int scsi_block_when_processing_errors (struct scsi_device *  
sdev);
```

## Arguments

*sdev*

Device on which we are performing recovery.

## Description

We block until the host is out of error recovery, and then check to see whether the host or the device is offline.

## Return value

0 when dev was taken offline by error recovery. 1 OK to proceed.

# scsi\_eh\_prep\_cmd

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_eh_prep_cmd` — Save a scsi command info as part of error recovery

## Synopsis

```
void scsi_eh_prep_cmd (struct scsi_cmd * scmd, struct  
scsi_eh_save * ses, unsigned char * cmd, int cmd_size,  
unsigned sense_bytes);
```

## Arguments

*scmd*

SCSI command structure to hijack

*ses*

structure to save restore information

*cmd*

CDB to send. Can be NULL if no new cmd is needed

*cmd\_size*

size in bytes of *cmd* (must be <= BLK\_MAX\_CDB)

*sense\_bytes*

size of sense data to copy. or 0 (if != 0 *cmd* is ignored)

## Description

This function is used to save a scsi command information before re-execution as part of the error recovery process. If *sense\_bytes* is 0 the command sent must be one that does not transfer any data. If *sense\_bytes* != 0 *cmd* is ignored and this functions sets up a REQUEST\_SENSE command and *cmd* buffers to read *sense\_bytes* into *scmd->sense\_buffer*.

# scsi\_eh\_restore\_cmd

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_eh_restore_cmd` — Restore a scsi command info as part of error recovery

## Synopsis

```
void scsi_eh_restore_cmd (struct scsi_cmd * scmd, struct
scsi_eh_save * ses);
```

## Arguments

*scmd*

SCSI command structure to restore

*ses*

saved information from a corresponding call to `scsi_eh_prep_cmd`

## Description

Undo any damage done by above `scsi_eh_prep_cmd`.

# scsi\_eh\_finish\_cmd

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_eh_finish_cmd` — Handle a cmd that eh is finished with.

## Synopsis

```
void scsi_eh_finish_cmd (struct scsi_cmnd * scmd, struct  
list_head * done_q);
```

## Arguments

*scmd*

Original SCSI cmd that eh has finished.

*done\_q*

Queue for processed commands.

## Notes

We don't want to use the normal command completion while we are still handling errors - it may cause other commands to be queued, and that would disturb what we are doing. Thus we really want to keep a list of pending commands for final completion, and once we are ready to leave error handling we handle completion for real.

# scsi\_eh\_get\_sense

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_eh_get_sense` — Get device sense data.

## Synopsis

```
int scsi_eh_get_sense (struct list_head * work_q, struct
list_head * done_q);
```

## Arguments

*work\_q*

Queue of commands to process.

*done\_q*

Queue of processed commands.

## Description

See if we need to request sense information. if so, then get it now, so we have a better idea of what to do.

## Notes

This has the unfortunate side effect that if a shost adapter does not automatically request sense information, we end up shutting it down before we request it.

All drivers should request sense information internally these days, so for now all I have to say is tough noogies if you end up in here.

## XXX

Long term this code should go away, but that needs an audit of all LLDDs first.

# scsi\_eh\_ready\_devs

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_eh_ready_devs` — check device ready state and recover if not.

## Synopsis

```
void scsi_eh_ready_devs (struct Scsi_Host * shost, struct  
list_head * work_q, struct list_head * done_q);
```

## Arguments

*shost*

host to be recovered.

*work\_q*

list\_head for pending commands.

*done\_q*

list\_head for processed commands.

## scsi\_eh\_flush\_done\_q

### LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_eh_flush_done_q` — finish processed commands or retry them.

### Synopsis

```
void scsi_eh_flush_done_q (struct list_head * done_q);
```

### Arguments

*done\_q*

list\_head of processed commands.

# scsi\_normalize\_sense

## LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_normalize_sense` — normalize main elements from either fixed or descriptor sense data format into a common format.

### Synopsis

```
int scsi_normalize_sense (const u8 * sense_buffer, int sb_len,  
struct scsi_sense_hdr * sshdr);
```

### Arguments

*sense\_buffer*

byte array containing sense data returned by device

*sb\_len*

number of valid bytes in *sense\_buffer*

*sshdr*

pointer to instance of structure that common elements are written to.

### Notes

The “main elements” from sense data are: `response_code`, `sense_key`, `asc`, `ascq` and `additional_length` (only for descriptor format).

Typically this function can be called after a device has responded to a SCSI command with the `CHECK_CONDITION` status.

## Return value

1 if valid sense data information found, else 0;

# scsi\_sense\_desc\_find

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_sense_desc_find` — search for a given descriptor type in descriptor sense data format.

## Synopsis

```
const u8 * scsi_sense_desc_find (const u8 * sense_buffer, int
sb_len, int desc_type);
```

## Arguments

*sense\_buffer*

byte array of descriptor format sense data

*sb\_len*

number of valid bytes in *sense\_buffer*

*desc\_type*

value of descriptor type to find (e.g. 0 -> information)

## Notes

only valid when sense data is in descriptor format

## Return value

pointer to start of (first) descriptor if found else NULL

# scsi\_get\_sense\_info\_fld

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_get_sense_info_fld` — get information field from sense data (either fixed or descriptor format)

## Synopsis

```
int scsi_get_sense_info_fld (const u8 * sense_buffer, int
sb_len, u64 * info_out);
```

## Arguments

*sense\_buffer*

byte array of sense data

*sb\_len*

number of valid bytes in *sense\_buffer*

*info\_out*

pointer to 64 integer where 8 or 4 byte information field will be placed if found.

## Return value

1 if information field found, 0 if not found.

# scsi\_build\_sense\_buffer

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_build_sense_buffer` — build sense data in a buffer

## Synopsis

```
void scsi_build_sense_buffer (int desc, u8 * buf, u8 key, u8  
asc, u8 ascq);
```

## Arguments

*desc*

Sense format (non zero == descriptor format, 0 == fixed format)

*buf*

Where to build sense data

*key*

Sense key

*asc*

Additional sense code

*ascq*

Additional sense code qualifier

### 3.1.5. drivers/scsi/scsi\_devinfo.c

Manage `scsi_dev_info_list`, which tracks blacklisted and whitelisted devices.

## scsi\_dev\_info\_list\_add

### LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_dev_info_list_add` — add one `dev_info` list entry.

### Synopsis

```
int scsi_dev_info_list_add (int compatible, char * vendor,  
char * model, char * strflags, int flags);
```

### Arguments

*compatible*

if true, null terminate short strings. Otherwise space pad.

*vendor*

vendor string

*model*

model (product) string

*strflags*

integer string

*flags*

if strflags NULL, use this flag value

## Description

Create and add one `dev_info` entry for *vendor*, *model*, *strflags* or *flag*. If *compatible*, add to the tail of the list, do not space pad, and set `devinfo->compatible`. The `scsi_static_device_list` entries are added with *compatible* 1 and *clflags* NULL.

## Returns

0 OK, -error on failure.

# scsi\_dev\_info\_list\_add\_str

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_dev_info_list_add_str` — parse `dev_list` and add to the `scsi_dev_info_list`.

## Synopsis

```
int scsi_dev_info_list_add_str (char * dev_list);
```

## Arguments

*dev\_list*

string of device flags to add

## Description

Parse *dev\_list*, and add entries to the *scsi\_dev\_info\_list*. *dev\_list* is of the form “vendor:product:flag,vendor:product:flag”. *dev\_list* is modified via *strsep*. Can be called for command line addition, for *proc* or maybe a *sysfs* interface.

## Returns

0 if OK, -error on failure.

# scsi\_get\_device\_flags

## LINUX

Kernel Hackers Manual July 2026

## Name

*scsi\_get\_device\_flags* — get device specific flags from the dynamic device list.

## Synopsis

```
int scsi_get_device_flags (struct scsi_device * sdev, const unsigned char * vendor, const unsigned char * model);
```

## Arguments

*sdev*

scsi\_device to get flags for

*vendor*

vendor name

*model*

model name

## Description

Search the global `scsi_dev_info_list` (specified by list zero) for an entry matching *vendor* and *model*, if found, return the matching flags value, else return the host or global default settings. Called during scan time.

## scsi\_exit\_devinfo

### LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_exit_devinfo` — remove `/proc/scsi/device_info` & the `scsi_dev_info_list`

### Synopsis

```
void scsi_exit_devinfo ( void );
```

## Arguments

*void*

no arguments

# scsi\_init\_devinfo

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_init_devinfo` — set up the dynamic device list.

## Synopsis

```
int scsi_init_devinfo ( void );
```

## Arguments

*void*

no arguments

## Description

Add command line entries from `scsi_dev_flags`, then add `scsi_static_device_list` entries to the scsi device info list.

## 3.1.6. drivers/scsi/scsi\_ioctl.c

Handle ioctl() calls for SCSI devices.

# scsi\_ioctl

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_ioctl` — Dispatch ioctl to scsi device

## Synopsis

```
int scsi_ioctl (struct scsi_device * sdev, int cmd, void  
__user * arg);
```

## Arguments

*sdev*

scsi device receiving ioctl

*cmd*

which ioctl is it

*arg*

data associated with ioctl

## Description

The `scsi_ioctl` function differs from most ioctls in that it does not take a major/minor number as the `dev` field. Rather, it takes a pointer to a struct `scsi_device`.

# scsi\_nonblockable\_ioctl

## LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_nonblockable_ioctl` — Handle SG\_SCSI\_RESET

### Synopsis

```
int scsi_nonblockable_ioctl (struct scsi_device * sdev, int cmd, void __user * arg, int ndelay);
```

### Arguments

*sdev*

scsi device receiving ioctl

*cmd*

Must be SC\_SCSI\_RESET

*arg*

pointer to int containing SG\_SCSI\_RESET\_{DEVICE,BUS,HOST}

*ndelay*

file mode O\_NDELAY flag

## 3.1.7. drivers/scsi/scsi\_lib.c

SCSI queuing library.

# scsi\_execute

## LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_execute` — insert request and wait for the result

### Synopsis

```
int scsi_execute (struct scsi_device * sdev, const unsigned
char * cmd, int data_direction, void * buffer, unsigned
bufflen, unsigned char * sense, int timeout, int retries, int
flags, int * resid);
```

### Arguments

*sdev*

scsi device

*cmd*

scsi command

*data\_direction*

data direction

*buffer*

data buffer

## Chapter 3. SCSI mid layer

*bufflen*

len of buffer

*sense*

optional sense buffer

*timeout*

request timeout in seconds

*retries*

number of times to retry request

*flags*

or into request flags;

*resid*

optional residual length

### Description

returns the req->errors value which is the scsi\_cmnd result field.

## scsi\_device\_from\_queue

**LINUX**

Kernel Hackers Manual July 2026

### Name

scsi\_device\_from\_queue — return sdev associated with a request\_queue

## Synopsis

```
struct scsi_device * scsi_device_from_queue (struct
request_queue * q);
```

## Arguments

*q*

The request queue to return the sdev from

## Description

Return the sdev associated with a request queue or NULL if the request\_queue does not reference a SCSI device.

# scsi\_mode\_select

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_mode_select` — issue a mode select

## Synopsis

```
int scsi_mode_select (struct scsi_device * sdev, int pf, int
sp, int modepage, unsigned char * buffer, int len, int
timeout, int retries, struct scsi_mode_data * data, struct
scsi_sense_hdr * sshdr);
```

## Arguments

*sdev*

SCSI device to be queried

*pf*

Page format bit (1 == standard, 0 == vendor specific)

*sp*

Save page bit (0 == don't save, 1 == save)

*modepage*

mode page being requested

*buffer*

request buffer (may not be smaller than eight bytes)

*len*

length of request buffer.

*timeout*

command timeout

*retries*

number of retries before failing

*data*

returns a structure abstracting the mode header data

*sshdr*

place to put sense data (or NULL if no sense to be collected). must be SCSI\_SENSE\_BUFFERSIZE big.

## Description

Returns zero if successful; negative error number or scsi status on error

# scsi\_mode\_sense

## LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_mode_sense` — issue a mode sense, falling back from 10 to six bytes if necessary.

### Synopsis

```
int scsi_mode_sense (struct scsi_device * sdev, int dbd, int
modepage, unsigned char * buffer, int len, int timeout, int
retries, struct scsi_mode_data * data, struct scsi_sense_hdr *
sshdr);
```

### Arguments

*sdev*

SCSI device to be queried

*dbd*

set if mode sense will allow block descriptors to be returned

*modepage*

mode page being requested

*buffer*

request buffer (may not be smaller than eight bytes)

*len*

length of request buffer.

*timeout*

command timeout

*retries*

number of retries before failing

*data*

returns a structure abstracting the mode header data

*sshdr*

place to put sense data (or NULL if no sense to be collected). must be SCSI\_SENSE\_BUFFERSIZE big.

## Description

Returns zero if unsuccessful, or the header offset (either 4 or 8 depending on whether a six or ten byte command was issued) if successful.

# scsi\_test\_unit\_ready

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_test_unit_ready` — test if unit is ready

## Synopsis

```
int scsi_test_unit_ready (struct scsi_device * sdev, int
timeout, int retries, struct scsi_sense_hdr * sshdr_external);
```

## Arguments

*sdev*

scsi device to change the state of.

*timeout*

command timeout

*retries*

number of retries before failing

*sshdr\_external*

Optional pointer to struct `scsi_sense_hdr` for returning sense. Make sure that this is cleared before passing in.

## Description

Returns zero if unsuccessful or an error if TUR failed. For removable media, `UNIT_ATTENTION` sets `->changed` flag.

# scsi\_device\_set\_state

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_device_set_state` — Take the given device through the device state model.

## Synopsis

```
int scsi_device_set_state (struct scsi_device * sdev, enum
scsi_device_state state);
```

## Arguments

*sdev*

scsi device to change the state of.

*state*

state to change to.

## Description

Returns zero if unsuccessful or an error if the requested transition is illegal.

# sdev\_evt\_send

## LINUX

Kernel Hackers Manual July 2026

## Name

`sdev_evt_send` — send asserted event to uevent thread

## Synopsis

```
void sdev_evt_send (struct scsi_device * sdev, struct  
scsi_event * evt);
```

## Arguments

*sdev*

scsi\_device event occurred on

*evt*

event to send

## Description

Assert scsi device event asynchronously.

# sdev\_evt\_alloc

## LINUX

Kernel Hackers Manual July 2026

## Name

`sdev_evt_alloc` — allocate a new scsi event

## Synopsis

```
struct scsi_event * sdev_evt_alloc (enum scsi_device_event  
evt_type, gfp_t gfpflags);
```

## Arguments

*evt\_type*

type of event to allocate

*gfpflags*

GFP flags for allocation

## Description

Allocates and returns a new `scsi_event`.

# sdev\_evt\_send\_simple

## LINUX

Kernel Hackers Manual July 2026

## Name

`sdev_evt_send_simple` — send asserted event to uevent thread

## Synopsis

```
void sdev_evt_send_simple (struct scsi_device * sdev, enum  
scsi_device_event evt_type, gfp_t gfpflags);
```

## Arguments

*sdev*

scsi\_device event occurred on

*evt\_type*

type of event to send

*gfpflags*

GFP flags for allocation

## Description

Assert scsi device event asynchronously, given an event type.

# scsi\_device\_quiesce

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_device_quiesce` — Block user issued commands.

## Synopsis

```
int scsi_device_quiesce (struct scsi_device * sdev);
```

## Arguments

*sdev*

scsi device to quiesce.

## Description

This works by trying to transition to the `SDEV_QUIESCE` state (which must be a legal transition). When the device is in this state, only special requests will be accepted, all others will be deferred. Since special requests may also be queued requests, a successful return doesn't guarantee the device will be totally quiescent.

Must be called with user context, may sleep.

Returns zero if unsuccessful or an error if not.

## scsi\_device\_resume

### LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_device_resume` — Restart user issued commands to a quiesced device.

### Synopsis

```
void scsi_device_resume (struct scsi_device * sdev);
```

### Arguments

*sdev*

scsi device to resume.

### Description

Moves the device from quiesced back to running and restarts the queues.

Must be called with user context, may sleep.

## scsi\_internal\_device\_block

### LINUX

## Name

`scsi_internal_device_block` — internal function to put a device temporarily into the SDEV\_BLOCK state

## Synopsis

```
int scsi_internal_device_block (struct scsi_device * sdev);
```

## Arguments

*sdev*

device to block

## Description

Block request made by scsi lld's to temporarily stop all scsi commands on the specified device. Called from interrupt or normal process context.

Returns zero if successful or error if not

## Notes

This routine transitions the device to the SDEV\_BLOCK state (which must be a legal transition). When the device is in this state, all commands are deferred until the scsi lld reenables the device with `scsi_device_unblock` or `device_block_tmo` fires. This routine assumes the `host_lock` is held on entry.

# scsi\_internal\_device\_unblock

## LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_internal_device_unblock` — resume a device after a block request

### Synopsis

```
int scsi_internal_device_unblock (struct scsi_device * sdev,  
enum scsi_device_state new_state);
```

### Arguments

*sdev*

device to resume

*new\_state*

state to set devices to after unblocking

### Description

Called by scsi lld's or the midlayer to restart the device queue for the previously suspended scsi device. Called from interrupt or normal process context.

Returns zero if successful or error if not.

### Notes

This routine transitions the device to the SDEV\_RUNNING state or to one of the offline states (which must be a legal transition) allowing the midlayer to goose the queue for this device. This routine assumes the `host_lock` is held upon entry.

# scsi\_kmap\_atomic\_sg

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_kmap_atomic_sg` — find and atomically map an sg-element

## Synopsis

```
void * scsi_kmap_atomic_sg (struct scatterlist * sgl, int
sg_count, size_t * offset, size_t * len);
```

## Arguments

*sgl*

scatter-gather list

*sg\_count*

number of segments in sg

*offset*

offset in bytes into sg, on return offset into the mapped area

*len*

bytes to map, on return number of bytes mapped

## Description

Returns virtual address of the start of the mapped page

## scsi\_kunmap\_atomic\_sg

### LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_kunmap_atomic_sg` — atomically unmap a virtual address, previously mapped with `scsi_kmap_atomic_sg`

### Synopsis

```
void scsi_kunmap_atomic_sg (void * virt);
```

### Arguments

*virt*

virtual address to be unmapped

### 3.1.8. drivers/scsi/scsi\_lib\_dma.c

SCSI library functions depending on DMA (map and unmap scatter-gather lists).

## scsi\_dma\_map

### LINUX

## Name

`scsi_dma_map` — perform DMA mapping against command's sg lists

## Synopsis

```
int scsi_dma_map (struct scsi_cmnd * cmd);
```

## Arguments

*cmd*

scsi command

## Description

Returns the number of sg lists actually used, zero if the sg lists is NULL, or -ENOMEM if the mapping failed.

# scsi\_dma\_unmap

## LINUX

## Name

`scsi_dma_unmap` — unmap command's sg lists mapped by `scsi_dma_map`

## Synopsis

```
void scsi_dma_unmap (struct scsi_cmnd * cmd);
```

## Arguments

*cmd*

scsi command

### 3.1.9. drivers/scsi/scsi\_module.c

The file `drivers/scsi/scsi_module.c` contains legacy support for old-style host templates. It should never be used by any new driver.

### 3.1.10. drivers/scsi/scsi\_proc.c

The functions in this file provide an interface between the PROC file system and the SCSI device drivers. It is mainly used for debugging, statistics and to pass information directly to the lowlevel driver. I.E. plumbing to manage `/proc/scsi/*`

## proc\_scsi\_read

### LINUX

Kernel Hackers Manual July 2026

### Name

`proc_scsi_read` — handle read from `/proc` by calling host's `proc_info` command

## Synopsis

```
int proc_scsi_read (char * buffer, char ** start, off_t
offset, int length, int * eof, void * data);
```

## Arguments

*buffer*

passed to `proc_info`

*start*

passed to `proc_info`

*offset*

passed to `proc_info`

*length*

passed to `proc_info`

*eof*

returns whether length read was less than requested

*data*

pointer to a struct `Scsi_Host`

## proc\_scsi\_write\_proc

### LINUX

Kernel Hackers Manual July 2026

### Name

`proc_scsi_write_proc` — Handle write to `/proc` by calling host's

proc\_info

## Synopsis

```
int proc_scsi_write_proc (struct file * file, const char  
__user * buf, unsigned long count, void * data);
```

## Arguments

*file*

not used

*buf*

source of data to write.

*count*

number of bytes (at most PROC\_BLOCK\_SIZE) to write.

*data*

pointer to struct Scsi\_Host

## scsi\_proc\_hostdir\_add

**LINUX**

Kernel Hackers Manual July 2026

### Name

scsi\_proc\_hostdir\_add — Create directory in /proc for a scsi host

## Synopsis

```
void scsi_proc_hostdir_add (struct scsi_host_template * sht);
```

## Arguments

*sht*

owner of this directory

## Description

Sets *sht->proc\_dir* to the new directory.

# scsi\_proc\_hostdir\_rm

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_proc_hostdir_rm` — remove directory in /proc for a scsi host

## Synopsis

```
void scsi_proc_hostdir_rm (struct scsi_host_template * sht);
```

## Arguments

*sh*t

owner of directory

## scsi\_proc\_host\_add

### LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_proc_host_add` — Add entry for this host to appropriate /proc dir

### Synopsis

```
void scsi_proc_host_add (struct Scsi_Host * shost);
```

### Arguments

*shost*

host to add

## scsi\_proc\_host\_rm

### LINUX

## Name

`scsi_proc_host_rm` — remove this host's entry from /proc

## Synopsis

```
void scsi_proc_host_rm (struct Scsi_Host * shost);
```

## Arguments

*shost*

which host

# proc\_print\_scsidevice

## LINUX

## Name

`proc_print_scsidevice` — return data about this host

## Synopsis

```
int proc_print_scsidevice (struct device * dev, void * data);
```

## Arguments

*dev*

A scsi device

*data*

struct seq\_file to output to.

## Description

prints Host, Channel, Id, Lun, Vendor, Model, Rev, Type, and revision.

# scsi\_add\_single\_device

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_add_single_device` — Respond to user request to probe for/add device

## Synopsis

```
int scsi_add_single_device (uint host, uint channel, uint id,  
uint lun);
```

## Arguments

*host*

user-supplied decimal integer

*channel*

user-supplied decimal integer

*id*

user-supplied decimal integer

*lun*

user-supplied decimal integer

## Description

called by writing “scsi add-single-device” to `/proc/scsi/scsi`.

does `scsi_host_lookup` and either `user_scan` if that transport type supports it, or else `scsi_scan_host_selected`

## Note

this seems to be aimed exclusively at SCSI parallel busses.

# scsi\_remove\_single\_device

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_remove_single_device` — Respond to user request to remove a device

## Synopsis

```
int scsi_remove_single_device (uint host, uint channel, uint
id, uint lun);
```

## Arguments

*host*

user-supplied decimal integer

*channel*

user-supplied decimal integer

*id*

user-supplied decimal integer

*lun*

user-supplied decimal integer

## Description

called by writing “scsi remove-single-device” to /proc/scsi/scsi. Does a `scsi_device_lookup` and `scsi_remove_device`

## proc\_scsi\_write

### LINUX

Kernel Hackers Manual July 2026

### Name

`proc_scsi_write` — handle writes to /proc/scsi/scsi

## Synopsis

```
ssize_t proc_scsi_write (struct file * file, const char __user
* buf, size_t length, loff_t * ppos);
```

## Arguments

*file*

not used

*buf*

buffer to write

*length*

length of buf, at most PAGE\_SIZE

*ppos*

not used

## Description

this provides a legacy mechanism to add or remove devices by Host, Channel, ID, and Lun. To use, “echo ‘scsi add-single-device 0 1 2 3’ > /proc/scsi/scsi” or “echo ‘scsi remove-single-device 0 1 2 3’ > /proc/scsi/scsi” with “0 1 2 3” replaced by the Host, Channel, Id, and Lun.

## Note

this seems to be aimed at parallel SCSI. Most modern busses (USB, SATA, Firewire, Fibre Channel, etc) dynamically assign these values to provide a unique identifier and nothing more.

# proc\_scsi\_open

## LINUX

Kernel Hackers Manual July 2026

### Name

`proc_scsi_open` — glue function

### Synopsis

```
int proc_scsi_open (struct inode * inode, struct file * file);
```

### Arguments

*inode*

not used

*file*

passed to `single_open`

### Description

Associates `proc_scsi_show` with this file

# scsi\_init\_procfs

## LINUX

## Name

`scsi_init_procfs` — create `scsi` and `scsi/scsi` in `procfs`

## Synopsis

```
int scsi_init_procfs ( void );
```

## Arguments

*void*

no arguments

# scsi\_exit\_procfs

## LINUX

## Name

`scsi_exit_procfs` — Remove `scsi/scsi` and `scsi` from `procfs`

## Synopsis

```
void scsi_exit_procfs ( void );
```

## Arguments

*void*

no arguments

### 3.1.11. drivers/scsi/scsi\_netlink.c

Infrastructure to provide async events from transports to userspace via netlink, using a single NETLINK\_SCSITRANSPORT protocol for all transports. See the original patch submission

(<http://marc.info/?l=linux-scsi&m=115507374832500&w=2>) for more details.

## scsi\_nl\_rcv\_msg

### LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_nl_rcv_msg` — Receive message handler.

### Synopsis

```
void scsi_nl_rcv_msg (struct sk_buff * skb);
```

### Arguments

*skb*

socket receive buffer

## Description

Extracts message from a receive buffer. Validates message header and calls appropriate transport message handler

# scsi\_nl\_rcv\_event

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_nl_rcv_event` — Event handler for a netlink socket.

## Synopsis

```
int scsi_nl_rcv_event (struct notifier_block * this, unsigned  
long event, void * ptr);
```

## Arguments

*this*

event notifier block

*event*

event type

*ptr*

event payload

# scsi\_generic\_msg\_handler

## LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_generic_msg_handler` — receive message handler for GENERIC transport messages

### Synopsis

```
int scsi_generic_msg_handler (struct sk_buff * skb);
```

### Arguments

*skb*

socket receive buffer

# scsi\_netlink\_init

## LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_netlink_init` — Called by SCSI subsystem to initialize the SCSI transport netlink interface

## Synopsis

```
void scsi_netlink_init ( void );
```

## Arguments

*void*

no arguments

## scsi\_netlink\_exit

### LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_netlink_exit` — Called by SCSI subsystem to disable the SCSI transport netlink interface

## Synopsis

```
void scsi_netlink_exit ( void );
```

## Arguments

*void*

no arguments

## Description

### 3.1.12. drivers/scsi/scsi\_scan.c

Scan a host to determine which (if any) devices are attached. The general scanning/probing algorithm is as follows, exceptions are made to it depending on device specific flags, compilation options, and global variable (boot or module load time) settings. A specific LUN is scanned via an INQUIRY command; if the LUN has a device attached, a `scsi_device` is allocated and setup for it. For every id of every channel on the given host, start by scanning LUN 0. Skip hosts that don't respond at all to a scan of LUN 0. Otherwise, if LUN 0 has a device attached, allocate and setup a `scsi_device` for it. If target is SCSI-3 or up, issue a REPORT LUN, and scan all of the LUNs returned by the REPORT LUN; else, sequentially scan LUNs up until some maximum is reached, or a LUN is seen that cannot have a device attached to it.

## scsi\_complete\_async\_scans

### LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_complete_async_scans` — Wait for asynchronous scans to complete

### Synopsis

```
int scsi_complete_async_scans ( void );
```

## Arguments

*void*

no arguments

## Description

When this function returns, any host which started scanning before this function was called will have finished its scan. Hosts which started scanning after this function was called may or may not have finished.

# scsi\_unlock\_floptical

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_unlock_floptical` — unlock device via a special MODE SENSE command

## Synopsis

```
void scsi_unlock_floptical (struct scsi_device * sdev,
unsigned char * result);
```

## Arguments

*sdev*

scsi device to send command to

*result*

area to store the result of the MODE SENSE

## Description

Send a vendor specific MODE SENSE (not a MODE SELECT) command. Called for BLIST\_KEY devices.

# scsi\_alloc\_sdev

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_alloc_sdev` — allocate and setup a `scsi_Device`

## Synopsis

```
struct scsi_device * scsi_alloc_sdev (struct scsi_target *  
target, unsigned int lun, void * hostdata);
```

## Arguments

*target*

which target to allocate a `scsi_device` for

*lun*

which lun

*hostdata*

usually NULL and set by `->slave_alloc` instead

## Description

Allocate, initialize for io, and return a pointer to a `scsi_Device`. Stores the *shost*, *channel*, *id*, and *lun* in the `scsi_Device`, and adds `scsi_Device` to the appropriate list.

## Return value

`scsi_Device` pointer, or NULL on failure.

# scsi\_alloc\_target

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_alloc_target` — allocate a new or find an existing target

## Synopsis

```
struct scsi_target * scsi_alloc_target (struct device *
parent, int channel, uint id);
```

## Arguments

*parent*

parent of the target (need not be a scsi host)

*channel*

target channel number (zero if no channels)

*id*

target id number

## Description

Return an existing target if one exists, provided it hasn't already gone into `STARGET_DEL` state, otherwise allocate a new target.

The target is returned with an incremented reference, so the caller is responsible for both reaping and doing a last put

## scsi\_target\_reap

### LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_target_reap` — check to see if target is in use and destroy if not

### Synopsis

```
void scsi_target_reap (struct scsi_target * target);
```

## Arguments

*starget*

target to be checked

## Description

This is used after removing a LUN or doing a last put of the target it checks atomically that nothing is using the target and removes it if so.

# sanitize\_inquiry\_string

## LINUX

Kernel Hackers Manual July 2026

## Name

`sanitize_inquiry_string` — remove non-graphical chars from an INQUIRY result string

## Synopsis

```
void sanitize_inquiry_string (unsigned char * s, int len);
```

## Arguments

*s*

INQUIRY result string to sanitize

*len*

length of the string

## Description

The SCSI spec says that INQUIRY vendor, product, and revision strings must consist entirely of graphic ASCII characters, padded on the right with spaces. Since not all devices obey this rule, we will replace non-graphic or non-ASCII characters with spaces. Exception: a NUL character is interpreted as a string terminator, so all the following characters are set to spaces.

# scsi\_probe\_lun

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_probe_lun` — probe a single LUN using a SCSI INQUIRY

## Synopsis

```
int scsi_probe_lun (struct scsi_device * sdev, unsigned char *  
inq_result, int result_len, int * bflags);
```

## Arguments

*sdev*

scsi\_device to probe

*inq\_result*

area to store the INQUIRY result

*result\_len*

len of *inq\_result*

*bflags*

store any bflags found here

## Description

Probe the lun associated with *req* using a standard SCSI INQUIRY;

If the INQUIRY is successful, zero is returned and the INQUIRY data is in *inq\_result*; the *scsi\_level* and INQUIRY length are copied to the *scsi\_device* any flags value is stored in *\*bflags*.

# scsi\_test\_lun

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_test_lun` — waiting for a LUN to become ready

## Synopsis

```
int scsi_test_lun (struct scsi_device * sdev);
```

## Arguments

*sdev*

scsi\_device to test

## Description

Wait for the lun associated with *sdev* to become ready

Send a TEST UNIT READY to detect any unit attention conditions. Retry TEST UNIT READY for up to *scsi\_scan\_timeout* if the returned sense key is 02/04/01 (Not ready, Logical Unit is in process of becoming ready)

## scsi\_add\_lun

### LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_add_lun` — allocate and fully initialize a `scsi_device`

### Synopsis

```
int scsi_add_lun (struct scsi_device * sdev, unsigned char *  
inq_result, int * bflags, int async);
```

### Arguments

*sdev*

holds information to be stored in the new `scsi_device`

*inq\_result*

holds the result of a previous INQUIRY to the LUN

*bflags*

black/white list flag

*async*

1 if this device is being scanned asynchronously

## Description

Initialize the `scsi_device` *sdev*. Optionally set fields based on values in *\*bflags*.

## SCSI\_SCAN\_NO\_RESPONSE

could not allocate or setup a `scsi_device`

## SCSI\_SCAN\_LUN\_PRESENT

a new `scsi_device` was allocated and initialized

# scsi\_inq\_str

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_inq_str` — print INQUIRY data from *min* to *max* index, strip trailing whitespace

## Synopsis

```
unsigned char * scsi_inq_str (unsigned char * buf, unsigned
char * inq, unsigned first, unsigned end);
```

## Arguments

*buf*

Output buffer with at least `end-first+1` bytes of space

*inq*

Inquiry buffer (input)

*first*

Offset of string into `inq`

*end*

Index after last character in `inq`

## scsi\_probe\_and\_add\_lun

### LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_probe_and_add_lun` — probe a LUN, if a LUN is found add it

### Synopsis

```
int scsi_probe_and_add_lun (struct scsi_target * starget, uint
lun, int * bflagsp, struct scsi_device ** sdevp, int rescan,
void * hostdata);
```

## Arguments

*starget*

pointer to target device structure

*lun*

LUN of target device

*bflagsp*

store bflags here if not NULL

*sdevp*

probe the LUN corresponding to this scsi\_device

*rescan*

if nonzero skip some code only needed on first scan

*hostdata*

passed to `scsi_alloc_sdev`

## Description

Call `scsi_probe_lun`, if a LUN with an attached device is found, allocate and set it up by calling `scsi_add_lun`.

### **SCSI\_SCAN\_NO\_RESPONSE**

could not allocate or setup a `scsi_device`

### **SCSI\_SCAN\_TARGET\_PRESENT**

target responded, but no device is attached at the LUN

### **SCSI\_SCAN\_LUN\_PRESENT**

a new `scsi_device` was allocated and initialized

# scsi\_sequential\_lun\_scan

## LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_sequential_lun_scan` — sequentially scan a SCSI target

### Synopsis

```
void scsi_sequential_lun_scan (struct scsi_target * target,  
int bflags, int scsi_level, int rescan);
```

### Arguments

*target*

pointer to target structure to scan

*bflags*

black/white list flag for LUN 0

*scsi\_level*

Which version of the standard does this device adhere to

*rescan*

passed to `scsi_probe_add_lun`

## Description

Generally, scan from LUN 1 (LUN 0 is assumed to already have been scanned) to some maximum lun until a LUN is found with no device attached. Use the bflags to figure out any oddities.

Modifies sdevscan->lun.

# scsi\_report\_lun\_scan

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_report_lun_scan` — Scan using SCSI REPORT LUN results

## Synopsis

```
int scsi_report_lun_scan (struct scsi_target * starget, int
bflags, int rescan);
```

## Arguments

*starget*

which target

*bflags*

Zero or a mix of BLIST\_NOLUN, BLIST\_REPORTLUN2, or BLIST\_NOREPORTLUN

*rescan*

nonzero if we can skip code only needed on first scan

## Description

Fast scanning for modern (SCSI-3) devices by sending a REPORT LUN command. Scan the resulting list of LUNs by calling `scsi_probe_and_add_lun`.

If `BLINK_REPORTLUN2` is set, scan a target that supports more than 8 LUNs even if it's older than SCSI-3. If `BLIST_NOREPORTLUN` is set, return 1 always. If `BLIST_NOLUN` is set, return 0 always. If `target->no_report_luns` is set, return 1 always.

### 0

scan completed (or no memory, so further scanning is futile)

### 1

could not scan with REPORT LUN

## `scsi_prep_async_scan`

### LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_prep_async_scan` — prepare for an async scan

### Synopsis

```
struct async_scan_data * scsi_prep_async_scan (struct  
Scsi_Host * shost);
```

## Arguments

*shost*

the host which will be scanned

## Returns

a cookie to be passed to `scsi_finish_async_scan`

Tells the midlayer this host is going to do an asynchronous scan. It reserves the host's position in the scanning list and ensures that other asynchronous scans started after this one won't affect the ordering of the discovered devices.

# scsi\_finish\_async\_scan

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_finish_async_scan` — asynchronous scan has finished

## Synopsis

```
void scsi_finish_async_scan (struct async_scan_data * data);
```

## Arguments

*data*

cookie returned from earlier call to `scsi_prep_async_scan`

## Description

All the devices currently attached to this host have been found. This function announces all the devices it has found to the rest of the system.

### 3.1.13. drivers/scsi/scsi\_sysctl.c

Set up the sysctl entry: `"/dev/scsi/logging_level"`  
(`DEV_SCSI_LOGGING_LEVEL`) which sets/returns `scsi_logging_level`.

### 3.1.14. drivers/scsi/scsi\_sysfs.c

SCSI sysfs interface routines.

## scsi\_remove\_device

### LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_remove_device` — unregister a device from the scsi bus

### Synopsis

```
void scsi_remove_device (struct scsi_device * sdev);
```

### Arguments

*sdev*

scsi\_device to unregister

# scsi\_remove\_target

## LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_remove_target` — try to remove a target and all its devices

### Synopsis

```
void scsi_remove_target (struct device * dev);
```

### Arguments

*dev*

generic starget or parent of generic stargets to be removed

### Note

This is slightly racy. It is possible that if the user requests the addition of another device then the target won't be removed.

## 3.1.15. drivers/scsi/hosts.c

mid to lowlevel SCSI driver interface

# scsi\_host\_set\_state

## LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_host_set_state` — Take the given host through the host state model.

### Synopsis

```
int scsi_host_set_state (struct Scsi_Host * shost, enum  
scsi_host_state state);
```

### Arguments

*shost*

scsi host to change the state of.

*state*

state to change to.

### Description

Returns zero if unsuccessful or an error if the requested transition is illegal.

# scsi\_remove\_host

## LINUX

## Name

`scsi_remove_host` — remove a scsi host

## Synopsis

```
void scsi_remove_host (struct Scsi_Host * shost);
```

## Arguments

*shost*

a pointer to a scsi host to remove

# scsi\_add\_host\_with\_dma

## LINUX

## Name

`scsi_add_host_with_dma` — add a scsi host with dma device

## Synopsis

```
int scsi_add_host_with_dma (struct Scsi_Host * shost, struct  
device * dev, struct device * dma_dev);
```

## Arguments

*shost*

scsi host pointer to add

*dev*

a struct device of type scsi class

*dma\_dev*

dma device for the host

## Note

You rarely need to worry about this unless you're in a virtualised host environments, so use the simpler `scsi_add_host` function instead.

## Return value

0 on success / != 0 for error

# scsi\_host\_alloc

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_host_alloc` — register a scsi host adapter instance.

## Synopsis

```
struct Scsi_Host * scsi_host_alloc (struct scsi_host_template  
* sht, int privsize);
```

## Arguments

*sht*

pointer to scsi host template

*privsize*

extra bytes to allocate for driver

## Note

Allocate a new Scsi\_Host and perform basic initialization. The host is not published to the scsi midlayer until `scsi_add_host` is called.

## Return value

Pointer to a new Scsi\_Host

# scsi\_host\_lookup

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_host_lookup` — get a reference to a Scsi\_Host by host no

## Synopsis

```
struct Scsi_Host * scsi_host_lookup (unsigned short hostnum);
```

## Arguments

*hostnum*

host number to locate

## Return value

A pointer to located Scsi\_Host or NULL.

The caller must do a `scsi_host_put` to drop the reference that `scsi_host_get` took. The `put_device` below dropped the reference from `class_find_device`.

# scsi\_host\_get

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_host_get` — inc a Scsi\_Host ref count

## Synopsis

```
struct Scsi_Host * scsi_host_get (struct Scsi_Host * shost);
```

## Arguments

*shost*

Pointer to Scsi\_Host to inc.

# scsi\_host\_put

## LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_host_put` — dec a Scsi\_Host ref count

### Synopsis

```
void scsi_host_put (struct Scsi_Host * shost);
```

### Arguments

*shost*

Pointer to Scsi\_Host to dec.

# scsi\_queue\_work

## LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_queue_work` — Queue work to the Scsi\_Host workqueue.

## Synopsis

```
int scsi_queue_work (struct Scsi_Host * shost, struct  
work_struct * work);
```

## Arguments

*shost*

Pointer to Scsi\_Host.

*work*

Work to queue for execution.

## Return value

1 - work queued for execution 0 - work is already queued -EINVAL - work queue doesn't exist

## scsi\_flush\_work

### LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_flush_work` — Flush a Scsi\_Host's workqueue.

## Synopsis

```
void scsi_flush_work (struct Scsi_Host * shost);
```

## Arguments

*shost*

Pointer to Scsi\_Host.

### 3.1.16. drivers/scsi/constants.c

mid to lowlevel SCSI driver interface

## scsi\_print\_status

### LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_print_status` — print scsi status description

### Synopsis

```
void scsi_print_status (unsigned char scsi_status);
```

### Arguments

*scsi\_status*

scsi status value

## Description

If the status is recognized, the description is printed. Otherwise “Unknown status” is output. No trailing space. If CONFIG\_SCSI\_CONSTANTS is not set, then print status in hex (e.g. “0x2” for Check Condition).

## 3.2. Transport classes

Transport classes are service libraries for drivers in the SCSI lower layer, which expose transport attributes in sysfs.

### 3.2.1. Fibre Channel transport

The file `drivers/scsi/scsi_transport_fc.c` defines transport attributes for Fibre Channel.

## fc\_get\_event\_number

### LINUX

Kernel Hackers Manual July 2026

### Name

`fc_get_event_number` — Obtain the next sequential FC event number

### Synopsis

```
u32 fc_get_event_number ( void );
```

## Arguments

*void*

no arguments

## Notes

We could have inlined this, but it would have required `fc_event_seq` to be exposed. For now, live with the subroutine call. Atomic used to avoid lock/unlock...

# fc\_host\_post\_event

## LINUX

Kernel Hackers Manual July 2026

## Name

`fc_host_post_event` — called to post an even on an `fc_host`.

## Synopsis

```
void fc_host_post_event (struct Scsi_Host * shost, u32  
event_number, enum fc_host_event_code event_code, u32  
event_data);
```

## Arguments

*shost*

host the event occurred on

## Chapter 3. SCSI mid layer

*event\_number*

fc event number obtained from `get_fc_event_number`

*event\_code*

fc\_host event being posted

*event\_data*

32bits of data for the event being posted

## Notes

This routine assumes no locks are held on entry.

# fc\_host\_post\_vendor\_event

## LINUX

Kernel Hackers Manual July 2026

## Name

`fc_host_post_vendor_event` — called to post a vendor unique event on an `fc_host`

## Synopsis

```
void fc_host_post_vendor_event (struct Scsi_Host * shost, u32  
event_number, u32 data_len, char * data_buf, u64 vendor_id);
```

## Arguments

*shost*

host the event occurred on

*event\_number*

fc event number obtained from `get_fc_event_number`

*data\_len*

amount, in bytes, of vendor unique data

*data\_buf*

pointer to vendor unique data

*vendor\_id*

Vendor id

## Notes

This routine assumes no locks are held on entry.

# fc\_remove\_host

## LINUX

Kernel Hackers Manual July 2026

## Name

`fc_remove_host` — called to terminate any `fc_transport`-related elements for a scsi host.

## Synopsis

```
void fc_remove_host (struct Scsi_Host * shost);
```

## Arguments

*shost*

Which Scsi\_Host

## Description

This routine is expected to be called immediately preceding the a driver's call to `scsi_remove_host`.

## WARNING

A driver utilizing the `fc_transport`, which fails to call this routine prior to `scsi_remove_host`, will leave dangling objects in `/sys/class/fc_remote_ports`. Access to any of these objects can result in a system crash !!!

## Notes

This routine assumes no locks are held on entry.

# fc\_remote\_port\_add

## LINUX

Kernel Hackers Manual July 2026

## Name

`fc_remote_port_add` — notify fc transport of the existence of a remote FC port.

## Synopsis

```
struct fc_rport * fc_remote_port_add (struct Scsi_Host *
shost, int channel, struct fc_rport_identifiers * ids);
```

## Arguments

*shost*

scsi host the remote port is connected to.

*channel*

Channel on shost port connected to.

*ids*

The world wide names, fc address, and FC4 port roles for the remote port.

## Description

The LLDD calls this routine to notify the transport of the existence of a remote port. The LLDD provides the unique identifiers (wwpn, wwn) of the port, its FC address (port\_id), and the FC4 roles that are active for the port.

For ports that are FCP targets (aka scsi targets), the FC transport maintains consistent target id bindings on behalf of the LLDD. A consistent target id binding is an assignment of a target id to a remote port identifier, which persists while the scsi host is attached. The remote port can disappear, then later reappear, and its target id assignment remains the same. This allows for shifts in FC addressing (if binding by wwpn or wwn) with no apparent changes to the scsi subsystem which is based on scsi host number and target id values. Bindings are only valid during the attachment of the scsi host. If the host detaches, then later re-attaches, target id bindings may change.

This routine is responsible for returning a remote port structure. The routine will search the list of remote ports it maintains internally on behalf of consistent target id mappings. If found, the remote port structure will be reused. Otherwise, a new remote port structure will be allocated.

Whenever a remote port is allocated, a new `fc_remote_port` class device is created.

Should not be called from interrupt context.

## Notes

This routine assumes no locks are held on entry.

# fc\_remote\_port\_delete

## LINUX

Kernel Hackers Manual July 2026

## Name

`fc_remote_port_delete` — notifies the fc transport that a remote port is no longer in existence.

## Synopsis

```
void fc_remote_port_delete (struct fc_rport * rport);
```

## Arguments

*rport*

The remote port that no longer exists

## Description

The LLDD calls this routine to notify the transport that a remote port is no longer part of the topology. Note: Although a port may no longer be part of the topology, it may persist in the remote ports displayed by the `fc_host`. We do this under 2 conditions: 1) If the port was a scsi target, we delay its deletion by “blocking” it. This allows the port to temporarily disappear, then reappear without disrupting the SCSI device tree attached to it. During the “blocked” period the port will still exist. 2) If the port was a scsi target and disappears for longer than we expect, we’ll delete

the port and the tear down the SCSI device tree attached to it. However, we want to semi-persist the target id assigned to that port if it eventually does exist. The port structure will remain (although with minimal information) so that the target id bindings remains.

If the remote port is not an FCP Target, it will be fully torn down and deallocated, including the `fc_remote_port` class device.

If the remote port is an FCP Target, the port will be placed in a temporary blocked state. From the LLDD's perspective, the rport no longer exists. From the SCSI midlayer's perspective, the SCSI target exists, but all sdevs on it are blocked from further I/O. The following is then expected.

If the remote port does not return (signaled by a LLDD call to `fc_remote_port_add`) within the `dev_loss_tmo` timeout, then the scsi target is removed - killing all outstanding i/o and removing the scsi devices attached to it. The port structure will be marked Not Present and be partially cleared, leaving only enough information to recognize the remote port relative to the scsi target id binding if it later appears. The port will remain as long as there is a valid binding (e.g. until the user changes the binding type or unloads the scsi host with the binding).

If the remote port returns within the `dev_loss_tmo` value (and matches according to the target id binding type), the port structure will be reused. If it is no longer a SCSI target, the target will be torn down. If it continues to be a SCSI target, then the target will be unblocked (allowing i/o to be resumed), and a scan will be activated to ensure that all luns are detected.

Called from normal process context only - cannot be called from interrupt.

## Notes

This routine assumes no locks are held on entry.

# fc\_remote\_port\_rolechg

**LINUX**

## Name

`fc_remote_port_rolechg` — notifies the fc transport that the roles on a remote may have changed.

## Synopsis

```
void fc_remote_port_rolechg (struct fc_rport * rport, u32  
roles);
```

## Arguments

*rport*

The remote port that changed.

*roles*

New roles for this port.

## Description

The LLDD calls this routine to notify the transport that the roles on a remote port may have changed. The largest effect of this is if a port now becomes a FCP Target, it must be allocated a scsi target id. If the port is no longer a FCP target, any scsi target id value assigned to it will persist in case the role changes back to include FCP Target. No changes in the scsi midlayer will be invoked if the role changes (in the expectation that the role will be resumed. If it doesn't normal error processing will take place).

Should not be called from interrupt context.

## Notes

This routine assumes no locks are held on entry.

# fc\_block\_scsi\_eh

## LINUX

Kernel Hackers Manual July 2026

### Name

`fc_block_scsi_eh` — Block SCSI eh thread for blocked `fc_rport`

### Synopsis

```
int fc_block_scsi_eh (struct scsi_cmnd * cmd);
```

### Arguments

*cmd*

SCSI command that `scsi_eh` is trying to recover

### Description

This routine can be called from a FC LLD `scsi_eh` callback. It blocks the `scsi_eh` thread until the `fc_rport` leaves the `FC_PORTSTATE_BLOCKED`, or the `fast_io_fail_tmo` fires. This is necessary to avoid the `scsi_eh` failing recovery actions for blocked `rports` which would lead to offlined SCSI devices.

### Returns

SUCCESS if the `fc_rport` left the state `FC_PORTSTATE_BLOCKED`.

FAST\_IO\_FAIL if the `fast_io_fail_tmo` fired, this should be passed back to `scsi_eh`.

# fc\_vport\_create

## LINUX

Kernel Hackers Manual July 2026

### Name

`fc_vport_create` — Admin App or LLDD requests creation of a vport

### Synopsis

```
struct fc_vport * fc_vport_create (struct Scsi_Host * shost,  
int channel, struct fc_vport_identifiers * ids);
```

### Arguments

*shost*

scsi host the virtual port is connected to.

*channel*

channel on shost port connected to.

*ids*

The world wide names, FC4 port roles, etc for the virtual port.

### Notes

This routine assumes no locks are held on entry.

# fc\_vport\_terminate

## LINUX

Kernel Hackers Manual July 2026

### Name

`fc_vport_terminate` — Admin App or LLDD requests termination of a vport

### Synopsis

```
int fc_vport_terminate (struct fc_vport * vport);
```

### Arguments

*vport*

fc\_vport to be terminated

### Description

Calls the LLDD `vport_delete` function, then deallocates and removes the vport from the shost and object tree.

### Notes

This routine assumes no locks are held on entry.

## 3.2.2. iSCSI transport class

The file `drivers/scsi/scsi_transport_iscsi.c` defines transport attributes for the iSCSI class, which sends SCSI packets over TCP/IP connections.

# iscsi\_create\_flashnode\_sess

## LINUX

Kernel Hackers Manual July 2026

### Name

`iscsi_create_flashnode_sess` — Add flashnode session entry in sysfs

### Synopsis

```
struct iscsi_bus_flash_session * iscsi_create_flashnode_sess
(struct Scsi_Host * shost, int index, struct iscsi_transport *
transport, int dd_size);
```

### Arguments

*shost*

pointer to host data

*index*

index of flashnode to add in sysfs

*transport*

pointer to transport data

*dd\_size*

total size to allocate

### Description

Adds a sysfs entry for the flashnode session attributes

## Returns

pointer to allocated flashnode sess on success NULL on failure

# iscsi\_create\_flashnode\_conn

## LINUX

Kernel Hackers Manual July 2026

## Name

`iscsi_create_flashnode_conn` — Add flashnode conn entry in sysfs

## Synopsis

```
struct iscsi_bus_flash_conn * iscsi_create_flashnode_conn
(struct Scsi_Host * shost, struct iscsi_bus_flash_session *
fnode_sess, struct iscsi_transport * transport, int dd_size);
```

## Arguments

*shost*

pointer to host data

*fnode\_sess*

pointer to the parent flashnode session entry

*transport*

pointer to transport data

*dd\_size*

total size to allocate

## Description

Adds a sysfs entry for the flashnode connection attributes

## Returns

pointer to allocated flashnode conn on success NULL on failure

# iscsi\_is\_flashnode\_conn\_dev

## LINUX

Kernel Hackers Manual July 2026

## Name

`iscsi_is_flashnode_conn_dev` — verify passed device is to be flashnode conn

## Synopsis

```
int iscsi_is_flashnode_conn_dev (struct device * dev, void * data);
```

## Arguments

*dev*

device to verify

*data*

pointer to data containing value to use for verification

## Description

Verifies if the passed device is flashnode conn device

## Returns

1 on success 0 on failure

# iscsi\_find\_flashnode\_sess

## LINUX

Kernel Hackers Manual July 2026

## Name

`iscsi_find_flashnode_sess` — finds flashnode session entry

## Synopsis

```
struct device * iscsi_find_flashnode_sess (struct Scsi_Host *  
shost, void * data, int (*fn) (struct device *dev, void  
*data));
```

## Arguments

*shost*

pointer to host data

*data*

pointer to data containing value to use for comparison

*fn*

function pointer that does actual comparison

## Description

Finds the flashnode session object comparing the data passed using logic defined in passed function pointer

## Returns

pointer to found flashnode session device object on success `NULL` on failure

# iscsi\_find\_flashnode\_conn

## LINUX

Kernel Hackers Manual July 2026

## Name

`iscsi_find_flashnode_conn` — finds flashnode connection entry

## Synopsis

```
struct device * iscsi_find_flashnode_conn (struct  
iscsi_bus_flash_session * fnode_sess, void * data, int (*fn)  
(struct device *dev, void *data));
```

## Arguments

*fnode\_sess*

pointer to parent flashnode session entry

*data*

pointer to data containing value to use for comparison

*fn*

function pointer that does actual comparison

## Description

Finds the flashnode connection object comparing the data passed using logic defined in passed function pointer

## Returns

pointer to found flashnode connection device object on success NULL on failure

# iscsi\_destroy\_flashnode\_sess

## LINUX

Kernel Hackers Manual July 2026

## Name

`iscsi_destroy_flashnode_sess` — destory flashnode session entry

## Synopsis

```
void iscsi_destroy_flashnode_sess (struct  
iscsi_bus_flash_session * fnode_sess);
```

## Arguments

*fnode\_sess*

pointer to flashnode session entry to be destroyed

## Description

Deletes the flashnode session entry and all children flashnode connection entries from sysfs

# iscsi\_destroy\_all\_flashnode

## LINUX

Kernel Hackers Manual July 2026

## Name

`iscsi_destroy_all_flashnode` — destory all flashnode session entries

## Synopsis

```
void iscsi_destroy_all_flashnode (struct Scsi_Host * shost);
```

## Arguments

*shost*

pointer to host data

## Description

Destroys all the flashnode session entries and all corresponding children flashnode connection entries from sysfs

# iscsi\_scan\_finished

## LINUX

Kernel Hackers Manual July 2026

## Name

`iscsi_scan_finished` — helper to report when running scans are done

## Synopsis

```
int iscsi_scan_finished (struct Scsi_Host * shost, unsigned long time);
```

## Arguments

*shost*

scsi host

*time*

scan run time

## Description

This function can be used by drives like `qla4xxx` to report to the scsi layer when the scans it kicked off at module load time are done.

## iscsi\_block\_scsi\_eh

### LINUX

Kernel Hackers Manual July 2026

### Name

`iscsi_block_scsi_eh` — block scsi eh until session state has transitioned

### Synopsis

```
int iscsi_block_scsi_eh (struct scsi_cmnd * cmd);
```

### Arguments

*cmd*

scsi cmd passed to scsi eh handler

### Description

If the session is down this function will wait for the recovery timer to fire or for the session to be logged back in. If the recovery timer fires then `FAST_IO_FAIL` is returned. The caller should pass this error value to the scsi eh.

## iscsi\_unblock\_session

### LINUX

## Name

`iscsi_unblock_session` — set a session as logged in and start IO.

## Synopsis

```
void iscsi_unblock_session (struct iscsi_cls_session *  
session);
```

## Arguments

*session*

iscsi session

## Description

Mark a session as ready to accept IO.

# iscsi\_create\_session

## LINUX

## Name

`iscsi_create_session` — create iscsi class session

## Synopsis

```
struct iscsi_cls_session * iscsi_create_session (struct  
Scsi_Host * shost, struct iscsi_transport * transport, int  
dd_size, unsigned int target_id);
```

## Arguments

*shost*

scsi host

*transport*

iscsi transport

*dd\_size*

private driver data size

*target\_id*

which target

## Description

This can be called from a LLD or `iscsi_transport`.

## `iscsi_destroy_session`

### LINUX

Kernel Hackers Manual July 2026

## Name

`iscsi_destroy_session` — destroy iscsi session

## Synopsis

```
int iscsi_destroy_session (struct iscsi_cls_session *
    session);
```

## Arguments

*session*

iscsi\_session

## Description

Can be called by a LLD or iscsi\_transport. There must not be any running connections.

## iscsi\_create\_conn

### LINUX

Kernel Hackers Manual July 2026

## Name

iscsi\_create\_conn — create iscsi class connection

## Synopsis

```
struct iscsi_cls_conn * iscsi_create_conn (struct
    iscsi_cls_session * session, int dd_size, uint32_t cid);
```

## Arguments

*session*

iscsi\_cls\_session

*dd\_size*

private driver data size

*cid*

connection id

## Description

This can be called from a LLD or `iscsi_transport`. The connection is child of the session so `cid` must be unique for all connections on the session.

Since we do not support MCS, `cid` will normally be zero. In some cases for software `iscsi` we could be trying to preallocate a connection struct in which case there could be two connection structs and `cid` would be non-zero.

## iscsi\_destroy\_conn

### LINUX

Kernel Hackers Manual July 2026

### Name

`iscsi_destroy_conn` — destroy iscsi class connection

### Synopsis

```
int iscsi_destroy_conn (struct iscsi_cls_conn * conn);
```

## Arguments

*conn*

iscsi\_cls\_session

## Description

This can be called from a LLD or iscsi\_transport.

# iscsi\_session\_event

## LINUX

Kernel Hackers Manual July 2026

## Name

`iscsi_session_event` — send session destr. completion event

## Synopsis

```
int iscsi_session_event (struct iscsi_cls_session * session,  
enum iscsi_uevent_e event);
```

## Arguments

*session*

iscsi class session

*event*

type of event

### 3.2.3. Serial Attached SCSI (SAS) transport class

The file `drivers/scsi/scsi_transport_sas.c` defines transport attributes for Serial Attached SCSI, a variant of SATA aimed at large high-end systems.

The SAS transport class contains common code to deal with SAS HBAs, an approximated representation of SAS topologies in the driver model, and various sysfs attributes to expose these topologies and management interfaces to userspace.

In addition to the basic SCSI core objects this transport class introduces two additional intermediate objects: The SAS PHY as represented by struct `sas_phy` defines an "outgoing" PHY on a SAS HBA or Expander, and the SAS remote PHY represented by struct `sas_rphy` defines an "incoming" PHY on a SAS Expander or end device. Note that this is purely a software concept, the underlying hardware for a PHY and a remote PHY is the exactly the same.

There is no concept of a SAS port in this code, users can see what PHYs form a wide port based on the `port_identifier` attribute, which is the same for all PHYs in a port.

## `sas_remove_children`

### LINUX

Kernel Hackers Manual July 2026

### Name

`sas_remove_children` — tear down a devices SAS data structures

### Synopsis

```
void sas_remove_children (struct device * dev);
```

## Arguments

*dev*

device belonging to the sas object

## Description

Removes all SAS PHYs and remote PHYs for a given object

# sas\_remove\_host

## LINUX

Kernel Hackers Manual July 2026

## Name

`sas_remove_host` — tear down a Scsi\_Host's SAS data structures

## Synopsis

```
void sas_remove_host (struct Scsi_Host * shost);
```

## Arguments

*shost*

Scsi Host that is torn down

## Description

Removes all SAS PHYs and remote PHYs for a given Scsi\_Host. Must be called just before `scsi_remove_host` for SAS HBAs.

# sas\_tlr\_supported

## LINUX

Kernel Hackers Manual July 2026

## Name

`sas_tlr_supported` — checking TLR bit in vpd 0x90

## Synopsis

```
unsigned int sas_tlr_supported (struct scsi_device * sdev);
```

## Arguments

*sdev*

scsi device struct

## Description

Check Transport Layer Retries are supported or not. If vpd page 0x90 is present, TLR is supported.

# **sas\_disable\_tlr**

## **LINUX**

Kernel Hackers Manual July 2026

### **Name**

`sas_disable_tlr` — setting TLR flags

### **Synopsis**

```
void sas_disable_tlr (struct scsi_device * sdev);
```

### **Arguments**

*sdev*

scsi device struct

### **Description**

Setting `tlr_enabled` flag to 0.

# **sas\_enable\_tlr**

## **LINUX**

Kernel Hackers Manual July 2026

### **Name**

`sas_enable_tlr` — setting TLR flags

## Synopsis

```
void sas_enable_tlr (struct scsi_device * sdev);
```

## Arguments

*sdev*

scsi device struct

## Description

Setting tlr\_enabled flag 1.

## sas\_phy\_alloc

### LINUX

Kernel Hackers Manual July 2026

## Name

`sas_phy_alloc` — allocates and initialize a SAS PHY structure

## Synopsis

```
struct sas_phy * sas_phy_alloc (struct device * parent, int  
number);
```

## Arguments

*parent*

Parent device

*number*

Phy index

## Description

Allocates an SAS PHY structure. It will be added in the device tree below the device specified by *parent*, which has to be either a `Scsi_Host` or `sas_rphy`.

## Returns

SAS PHY allocated or `NULL` if the allocation failed.

# sas\_phy\_add

## LINUX

Kernel Hackers Manual July 2026

## Name

`sas_phy_add` — add a SAS PHY to the device hierarchy

## Synopsis

```
int sas_phy_add (struct sas_phy * phy);
```

## Arguments

*phy*

The PHY to be added

## Description

Publishes a SAS PHY to the rest of the system.

# **sas\_phy\_free**

## **LINUX**

Kernel Hackers Manual July 2026

## **Name**

`sas_phy_free` — free a SAS PHY

## **Synopsis**

```
void sas_phy_free (struct sas_phy * phy);
```

## **Arguments**

*phy*

SAS PHY to free

## **Description**

Frees the specified SAS PHY.

## Note

This function must only be called on a PHY that has not successfully been added using `sas_phy_add`.

# sas\_phy\_delete

## LINUX

Kernel Hackers Manual July 2026

## Name

`sas_phy_delete` — remove SAS PHY

## Synopsis

```
void sas_phy_delete (struct sas_phy * phy);
```

## Arguments

*phy*

SAS PHY to remove

## Description

Removes the specified SAS PHY. If the SAS PHY has an associated remote PHY it is removed before.

## scsi\_is\_sas\_phy

### LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_is_sas_phy` — check if a struct device represents a SAS PHY

### Synopsis

```
int scsi_is_sas_phy (const struct device * dev);
```

### Arguments

*dev*

device to check

### Returns

1 if the device represents a SAS PHY, 0 else

## sas\_port\_add

### LINUX

Kernel Hackers Manual July 2026

### Name

`sas_port_add` — add a SAS port to the device hierarchy

## Synopsis

```
int sas_port_add (struct sas_port * port);
```

## Arguments

*port*

port to be added

## Description

publishes a port to the rest of the system

## **sas\_port\_free**

### **LINUX**

Kernel Hackers Manual July 2026

## **Name**

`sas_port_free` — free a SAS PORT

## **Synopsis**

```
void sas_port_free (struct sas_port * port);
```

## Arguments

*port*

SAS PORT to free

## Description

Frees the specified SAS PORT.

## Note

This function must only be called on a PORT that has not successfully been added using `sas_port_add`.

# sas\_port\_delete

## LINUX

Kernel Hackers Manual July 2026

## Name

`sas_port_delete` — remove SAS PORT

## Synopsis

```
void sas_port_delete (struct sas_port * port);
```

## Arguments

*port*

SAS PORT to remove

## Description

Removes the specified SAS PORT. If the SAS PORT has an associated phys, unlink them from the port as well.

# scsi\_is\_sas\_port

## LINUX

Kernel Hackers Manual July 2026

## Name

`scsi_is_sas_port` — check if a struct device represents a SAS port

## Synopsis

```
int scsi_is_sas_port (const struct device * dev);
```

## Arguments

*dev*

device to check

## Returns

1 if the device represents a SAS Port, 0 else

# sas\_port\_get\_phy

## LINUX

Kernel Hackers Manual July 2026

## Name

`sas_port_get_phy` — try to take a reference on a port member

## Synopsis

```
struct sas_phy * sas_port_get_phy (struct sas_port * port);
```

## Arguments

*port*

port to check

# sas\_port\_add\_phy

## LINUX

## Name

`sas_port_add_phy` — add another phy to a port to form a wide port

## Synopsis

```
void sas_port_add_phy (struct sas_port * port, struct sas_phy  
* phy);
```

## Arguments

*port*

port to add the phy to

*phy*

phy to add

## Description

When a port is initially created, it is empty (has no phys). All ports must have at least one phy to operated, and all wide ports must have at least two. The current code makes no difference between ports and wide ports, but the only object that can be connected to a remote device is a port, so ports must be formed on all devices with phys if they're connected to anything.

## `sas_port_delete_phy`

**LINUX**

## Name

`sas_port_delete_phy` — remove a phy from a port or wide port

## Synopsis

```
void sas_port_delete_phy (struct sas_port * port, struct  
sas_phy * phy);
```

## Arguments

*port*

port to remove the phy from

*phy*

phy to remove

## Description

This operation is used for tearing down ports again. It must be done to every port or wide port before calling `sas_port_delete`.

# `sas_end_device_alloc`

**LINUX**

## Name

`sas_end_device_alloc` — allocate an rphy for an end device

## Synopsis

```
struct sas_rphy * sas_end_device_alloc (struct sas_port *  
parent);
```

## Arguments

*parent*

which port

## Description

Allocates an SAS remote PHY structure, connected to *parent*.

## Returns

SAS PHY allocated or `NULL` if the allocation failed.

# `sas_expander_alloc`

**LINUX**

## Name

`sas_expander_alloc` — allocate an rphy for an end device

## Synopsis

```
struct sas_rphy * sas_expander_alloc (struct sas_port *  
parent, enum sas_device_type type);
```

## Arguments

*parent*

which port

*type*

SAS\_EDGE\_EXPANDER\_DEVICE or  
SAS\_FANOUT\_EXPANDER\_DEVICE

## Description

Allocates an SAS remote PHY structure, connected to *parent*.

## Returns

SAS PHY allocated or `NULL` if the allocation failed.

## `sas_rphy_add`

**LINUX**

## Name

`sas_rphy_add` — add a SAS remote PHY to the device hierarchy

## Synopsis

```
int sas_rphy_add (struct sas_rphy * rphy);
```

## Arguments

*rphy*

The remote PHY to be added

## Description

Publishes a SAS remote PHY to the rest of the system.

# `sas_rphy_free`

## LINUX

## Name

`sas_rphy_free` — free a SAS remote PHY

## Synopsis

```
void sas_rphy_free (struct sas_rphy * rphy);
```

## Arguments

*rphy*

SAS remote PHY to free

## Description

Frees the specified SAS remote PHY.

## Note

This function must only be called on a remote PHY that has not successfully been added using `sas_rphy_add` (or has been `sas_rphy_remove'd`)

# **sas\_rphy\_delete**

## **LINUX**

Kernel Hackers Manual July 2026

## **Name**

`sas_rphy_delete` — remove and free SAS remote PHY

## **Synopsis**

```
void sas_rphy_delete (struct sas_rphy * rphy);
```

## Arguments

*rphy*

SAS remote PHY to remove and free

## Description

Removes the specified SAS remote PHY and frees it.

# sas\_rphy\_unlink

## LINUX

Kernel Hackers Manual July 2026

## Name

`sas_rphy_unlink` — unlink SAS remote PHY

## Synopsis

```
void sas_rphy_unlink (struct sas_rphy * rphy);
```

## Arguments

*rphy*

SAS remote phy to unlink from its parent port

## Description

Removes port reference to an rphy

# sas\_rphy\_remove

## LINUX

Kernel Hackers Manual July 2026

## Name

`sas_rphy_remove` — remove SAS remote PHY

## Synopsis

```
void sas_rphy_remove (struct sas_rphy * rphy);
```

## Arguments

*rphy*

SAS remote phy to remove

## Description

Removes the specified SAS remote PHY.

# scsi\_is\_sas\_rphy

## LINUX

Kernel Hackers Manual July 2026

### Name

`scsi_is_sas_rphy` — check if a struct device represents a SAS remote PHY

### Synopsis

```
int scsi_is_sas_rphy (const struct device * dev);
```

### Arguments

*dev*

device to check

### Returns

1 if the device represents a SAS remote PHY, 0 else

# sas\_attach\_transport

## LINUX

Kernel Hackers Manual July 2026

### Name

`sas_attach_transport` — instantiate SAS transport template

## Synopsis

```
struct scsi_transport_template * sas_attach_transport (struct  
sas_function_template * ft);
```

## Arguments

*ft*

SAS transport class function template

## **sas\_release\_transport**

### **LINUX**

Kernel Hackers Manual July 2026

## Name

`sas_release_transport` — release SAS transport template instance

## Synopsis

```
void sas_release_transport (struct scsi_transport_template *  
t);
```

## Arguments

*t*

transport template instance

### 3.2.4. SATA transport class

The SATA transport is handled by libata, which has its own book of documentation in this directory.

### 3.2.5. Parallel SCSI (SPI) transport class

The file `drivers/scsi/scsi_transport_spi.c` defines transport attributes for traditional (fast/wide/ultra) SCSI busses.

## spi\_schedule\_dv\_device

### LINUX

Kernel Hackers Manual July 2026

### Name

`spi_schedule_dv_device` — schedule domain validation to occur on the device

### Synopsis

```
void spi_schedule_dv_device (struct scsi_device * sdev);
```

### Arguments

*sdev*

The device to validate

## Description

Identical to `spi_dv_device` above, except that the DV will be scheduled to occur in a workqueue later. All memory allocations are atomic, so may be called from any context including those holding SCSI locks.

# spi\_display\_xfer\_agreement

## LINUX

Kernel Hackers Manual July 2026

## Name

`spi_display_xfer_agreement` — Print the current target transfer agreement

## Synopsis

```
void spi_display_xfer_agreement (struct scsi_target *  
starget);
```

## Arguments

*starget*

The target for which to display the agreement

## Description

Each SPI port is required to maintain a transfer agreement for each other port on the bus. This function prints a one-line summary of the current agreement; more detailed information is available in `sysfs`.

### 3.2.6. SCSI RDMA (SRP) transport class

The file `drivers/scsi/scsi_transport_srp.c` defines transport attributes for SCSI over Remote Direct Memory Access.

## srp\_rport\_add

### LINUX

Kernel Hackers Manual July 2026

### Name

`srp_rport_add` — add a SRP remote port to the device hierarchy

### Synopsis

```
struct srp_rport * srp_rport_add (struct Scsi_Host * shost,
struct srp_rport_identifiers * ids);
```

### Arguments

*shost*

scsi host the remote port is connected to.

*ids*

The port id for the remote port.

### Description

Publishes a port to the rest of the system.

## srp\_rport\_del

### LINUX

Kernel Hackers Manual July 2026

### Name

`srp_rport_del` — remove a SRP remote port

### Synopsis

```
void srp_rport_del (struct srp_rport * rport);
```

### Arguments

*rport*

SRP remote port to remove

### Description

Removes the specified SRP remote port.

## srp\_remove\_host

### LINUX

## Name

`srp_remove_host` — tear down a `Scsi_Host`'s SRP data structures

## Synopsis

```
void srp_remove_host (struct Scsi_Host * shost);
```

## Arguments

*shost*

Scsi Host that is torn down

## Description

Removes all SRP remote ports for a given `Scsi_Host`. Must be called just before `scsi_remove_host` for SRP HBAs.

# srp\_attach\_transport

## LINUX

## Name

`srp_attach_transport` — instantiate SRP transport template

## Synopsis

```
struct scsi_transport_template * srp_attach_transport (struct  
srp_function_template * ft);
```

## Arguments

*ft*

SRP transport class function template

## srp\_release\_transport

### LINUX

Kernel Hackers Manual July 2026

## Name

`srp_release_transport` — release SRP transport template instance

## Synopsis

```
void srp_release_transport (struct scsi_transport_template *  
t);
```

## Arguments

*t*

transport template instance





# Chapter 4. SCSI lower layer

## 4.1. Host Bus Adapter transport types

Many modern device controllers use the SCSI command set as a protocol to communicate with their devices through many different types of physical connections.

In SCSI language a bus capable of carrying SCSI commands is called a "transport", and a controller connecting to such a bus is called a "host bus adapter" (HBA).

### 4.1.1. Debug transport

The file `drivers/scsi/scsi_debug.c` simulates a host adapter with a variable number of disks (or disk like devices) attached, sharing a common amount of RAM. Does a lot of checking to make sure that we are not getting blocks mixed up, and panics the kernel if anything out of the ordinary is seen.

To be more realistic, the simulated devices have the transport attributes of SAS disks.

For documentation see <http://sg.danny.cz/sg/sdebug26.html>

### 4.1.2. todo

Parallel (fast/wide/ultra) SCSI, USB, SATA, SAS, Fibre Channel, FireWire, ATAPI devices, Infiniband, I20, iSCSI, Parallel ports, netlink...

