

**NAME**

**pci**, **pci\_alloc\_msi**, **pci\_alloc\_msix**, **pci\_disable\_busmaster**, **pci\_disable\_io**, **pci\_enable\_busmaster**, **pci\_enable\_io**, **pci\_find\_bsf**, **pci\_find\_cap**, **pci\_find\_dbsf**, **pci\_find\_device**, **pci\_find\_extcap**, **pci\_find\_htcap**, **pci\_find\_next\_cap**, **pci\_find\_next\_extcap**, **pci\_find\_next\_htcap**, **pci\_find\_pcie\_root\_port**, **pci\_get\_id**, **pci\_get\_max\_payload**, **pci\_get\_max\_read\_req**, **pci\_get\_powerstate**, **pci\_get\_vpd\_ident**, **pci\_get\_vpd\_readonly**, **pci\_iov\_attach**, **pci\_iov\_attach\_name**, **pci\_iov\_detach**, **pci\_msi\_count**, **pci\_msix\_count**, **pci\_msix\_pba\_bar**, **pci\_msix\_table\_bar**, **pci\_pending\_msix**, **pci\_read\_config**, **pci\_release\_msi**, **pci\_remap\_msix**, **pci\_restore\_state**, **pci\_save\_state**, **pci\_set\_max\_read\_req**, **pci\_set\_powerstate**, **pci\_write\_config**, **pcie\_adjust\_config**, **pcie\_flr**, **pcie\_get\_max\_completion\_timeout**, **pcie\_read\_config**, **pcie\_wait\_for\_pending\_transactions**, **pcie\_write\_config** - PCI bus interface

**SYNOPSIS**

```
#include <sys/bus.h>
#include <dev/pci/pci.h>
#include <dev/pci/pcivar.h>
```

*int*

```
pci_alloc_msi(device_t dev, int *count);
```

*int*

```
pci_alloc_msix(device_t dev, int *count);
```

*int*

```
pci_disable_busmaster(device_t dev);
```

*int*

```
pci_disable_io(device_t dev, int space);
```

*int*

```
pci_enable_busmaster(device_t dev);
```

*int*

```
pci_enable_io(device_t dev, int space);
```

*device\_t*

```
pci_find_bsf(uint8_t bus, uint8_t slot, uint8_t func);
```

*int*

```
pci_find_cap(device_t dev, int capability, int *capreg);
```

*device\_t*

**pci\_find\_dbstf**(*uint32\_t domain, uint8\_t bus, uint8\_t slot, uint8\_t func*);

*device\_t*

**pci\_find\_device**(*uint16\_t vendor, uint16\_t device*);

*int*

**pci\_find\_extcap**(*device\_t dev, int capability, int \*capreg*);

*int*

**pci\_find\_htcap**(*device\_t dev, int capability, int \*capreg*);

*int*

**pci\_find\_next\_cap**(*device\_t dev, int capability, int start, int \*capreg*);

*int*

**pci\_find\_next\_extcap**(*device\_t dev, int capability, int start, int \*capreg*);

*int*

**pci\_find\_next\_htcap**(*device\_t dev, int capability, int start, int \*capreg*);

*device\_t*

**pci\_find\_pcie\_root\_port**(*device\_t dev*);

*int*

**pci\_get\_id**(*device\_t dev, enum pci\_id\_type type, uintptr\_t \*id*);

*int*

**pci\_get\_max\_payload**(*device\_t dev*);

*int*

**pci\_get\_max\_read\_req**(*device\_t dev*);

*int*

**pci\_get\_powerstate**(*device\_t dev*);

*int*

**pci\_get\_vpd\_ident**(*device\_t dev, const char \*\*identptr*);

*int*

**pci\_get\_vpd\_readonly**(*device\_t dev, const char \*kw, const char \*\*vptr*);

*int*

**pci\_msi\_count**(*device\_t dev*);

*int*

**pci\_msix\_count**(*device\_t dev*);

*int*

**pci\_msix\_pba\_bar**(*device\_t dev*);

*int*

**pci\_msix\_table\_bar**(*device\_t dev*);

*int*

**pci\_pending\_msix**(*device\_t dev, u\_int index*);

*uint32\_t*

**pci\_read\_config**(*device\_t dev, int reg, int width*);

*int*

**pci\_release\_msi**(*device\_t dev*);

*int*

**pci\_remap\_msix**(*device\_t dev, int count, const u\_int \*vectors*);

*void*

**pci\_restore\_state**(*device\_t dev*);

*void*

**pci\_save\_state**(*device\_t dev*);

*int*

**pci\_set\_max\_read\_req**(*device\_t dev, int size*);

*int*

**pci\_set\_powerstate**(*device\_t dev, int state*);

*void*

**pci\_write\_config**(*device\_t dev, int reg, uint32\_t val, int width*);

*uint32\_t*

**pcie\_adjust\_config**(*device\_t dev, int reg, uint32\_t mask, uint32\_t val, int width*);

*bool*

**pcie\_flr**(*device\_t dev, u\_int max\_delay, bool force*);

*int*

**pcie\_get\_max\_completion\_timeout**(*device\_t dev*);

*uint32\_t*

**pcie\_read\_config**(*device\_t dev, int reg, int width*);

*bool*

**pcie\_wait\_for\_pending\_transactions**(*device\_t dev, u\_int max\_delay*);

*void*

**pcie\_write\_config**(*device\_t dev, int reg, uint32\_t val, int width*);

*void*

**pci\_event\_fn**(*void \*arg, device\_t dev*);

**EVENTHANDLER\_REGISTER**(*pci\_add\_device, pci\_event\_fn*);

**EVENTHANDLER\_DEREGISTER**(*pci\_delete\_resource, pci\_event\_fn*);

**#include** <dev/pci/pci\_iov.h>

*int*

**pci\_iov\_attach**(*device\_t dev, nvlist\_t \*pf\_schema, nvlist\_t \*vf\_schema*);

*int*

**pci\_iov\_attach\_name**(*device\_t dev, nvlist\_t \*pf\_schema, nvlist\_t \*vf\_schema, const char \*fmt, ...*);

*int*

**pci\_iov\_detach**(*device\_t dev*);

## DESCRIPTION

The **pci** set of functions are used for managing PCI devices. The functions are split into several groups: raw configuration access, locating devices, device information, device configuration, and message signaled interrupts.

### Raw Configuration Access

The **pci\_read\_config()** function is used to read data from the PCI configuration space of the device *dev*, at offset *reg*, with *width* specifying the size of the access.

The **pci\_write\_config()** function is used to write the value *val* to the PCI configuration space of the device *dev*, at offset *reg*, with *width* specifying the size of the access.

The **pcie\_adjust\_config()** function is used to modify the value of a register in the PCI-express capability register set of device *dev*. The offset *reg* specifies a relative offset in the register set with *width* specifying the size of the access. The new value of the register is computed by modifying bits set in *mask* to the value in *val*. Any bits not specified in *mask* are preserved. The previous value of the register is returned.

The **pcie\_read\_config()** function is used to read the value of a register in the PCI-express capability register set of device *dev*. The offset *reg* specifies a relative offset in the register set with *width* specifying the size of the access.

The **pcie\_write\_config()** function is used to write the value *val* to a register in the PCI-express capability register set of device *dev*. The offset *reg* specifies a relative offset in the register set with *width* specifying the size of the access.

**NOTE:** Device drivers should only use these functions for functionality that is not available via another **pci()** function.

### Locating Devices

The **pci\_find\_bsf()** function looks up the *device\_t* of a PCI device, given its *bus*, *slot*, and *func*. The *slot* number actually refers to the number of the device on the bus, which does not necessarily indicate its geographic location in terms of a physical slot. Note that in case the system has multiple PCI domains, the **pci\_find\_bsf()** function only searches the first one. Actually, it is equivalent to:

```
pci_find_dbsf(0, bus, slot, func);
```

The **pci\_find\_dbsf()** function looks up the *device\_t* of a PCI device, given its *domain*, *bus*, *slot*, and *func*. The *slot* number actually refers to the number of the device on the bus, which does not necessarily indicate its geographic location in terms of a physical slot.

The **pci\_find\_device()** function looks up the *device\_t* of a PCI device, given its *vendor* and *device* IDs. Note that there can be multiple matches for this search; this function only returns the first matching device.

## Device Information

The **pci\_find\_cap()** function is used to locate the first instance of a PCI capability register set for the device *dev*. The capability to locate is specified by ID via *capability*. Constant macros of the form **PCIY\_XXX** for standard capability IDs are defined in *<dev/pci/pci.h>*. If the capability is found, then *\*capreg* is set to the offset in configuration space of the capability register set, and **pci\_find\_cap()** returns zero. If the capability is not found or the device does not support capabilities, **pci\_find\_cap()** returns an error. The **pci\_find\_next\_cap()** function is used to locate the next instance of a PCI capability register set for the device *dev*. The *start* should be the *\*capreg* returned by a prior **pci\_find\_cap()** or **pci\_find\_next\_cap()**. When no more instances are located **pci\_find\_next\_cap()** returns an error.

The **pci\_find\_extcap()** function is used to locate the first instance of a PCI-express extended capability register set for the device *dev*. The extended capability to locate is specified by ID via *capability*. Constant macros of the form **PCIZ\_XXX** for standard extended capability IDs are defined in *<dev/pci/pci.h>*. If the extended capability is found, then *\*capreg* is set to the offset in configuration space of the extended capability register set, and **pci\_find\_extcap()** returns zero. If the extended capability is not found or the device is not a PCI-express device, **pci\_find\_extcap()** returns an error. The **pci\_find\_next\_extcap()** function is used to locate the next instance of a PCI-express extended capability register set for the device *dev*. The *start* should be the *\*capreg* returned by a prior **pci\_find\_extcap()** or **pci\_find\_next\_extcap()**. When no more instances are located **pci\_find\_next\_extcap()** returns an error.

The **pci\_find\_htcap()** function is used to locate the first instance of a HyperTransport capability register set for the device *dev*. The capability to locate is specified by type via *capability*. Constant macros of the form **PCIM\_HTCAP\_XXX** for standard HyperTransport capability types are defined in *<dev/pci/pci.h>*. If the capability is found, then *\*capreg* is set to the offset in configuration space of the capability register set, and **pci\_find\_htcap()** returns zero. If the capability is not found or the device is not a HyperTransport device, **pci\_find\_htcap()** returns an error. The **pci\_find\_next\_htcap()** function is used to locate the next instance of a HyperTransport capability register set for the device *dev*. The *start* should be the *\*capreg* returned by a prior **pci\_find\_htcap()** or **pci\_find\_next\_htcap()**. When no more instances are located **pci\_find\_next\_htcap()** returns an error.

The **pci\_find\_pcie\_root\_port()** function walks up the PCI device hierarchy to locate the PCI-express root port upstream of *dev*. If a root port is not found, **pci\_find\_pcie\_root\_port()** returns NULL.

The **pci\_get\_id()** function is used to read an identifier from a device. The *type* flag is used to specify which identifier to read. The following flags are supported:

**PCI\_ID\_RID** Read the routing identifier for the device.

**PCI\_ID\_MSI** Read the MSI routing ID. This is needed by some interrupt controllers to route MSI and MSI-X interrupts.

The **pci\_get\_vpd\_ident()** function is used to fetch a device's Vital Product Data (VPD) identifier string. If the device *dev* supports VPD and provides an identifier string, then *\*identptr* is set to point at a read-only, null-terminated copy of the identifier string, and **pci\_get\_vpd\_ident()** returns zero. If the device does not support VPD or does not provide an identifier string, then **pci\_get\_vpd\_ident()** returns an error.

The **pci\_get\_vpd\_readonly()** function is used to fetch the value of a single VPD read-only keyword for the device *dev*. The keyword to fetch is identified by the two character string *kw*. If the device supports VPD and provides a read-only value for the requested keyword, then *\*vptr* is set to point at a read-only, null-terminated copy of the value, and **pci\_get\_vpd\_readonly()** returns zero. If the device does not support VPD or does not provide the requested keyword, then **pci\_get\_vpd\_readonly()** returns an error.

The **pcie\_get\_max\_completion\_timeout()** function returns the maximum completion timeout configured for the device *dev* in microseconds. If the *dev* device is not a PCI-express device, **pcie\_get\_max\_completion\_timeout()** returns zero. When completion timeouts are disabled for *dev*, this function returns the maximum timeout that would be used if timeouts were enabled.

The **pcie\_wait\_for\_pending\_transactions()** function waits for any pending transactions initiated by the *dev* device to complete. The function checks for pending transactions by polling the transactions pending flag in the PCI-express device status register. It returns true once the transaction pending flag is clear. If transactions are still pending after *max\_delay* milliseconds, **pcie\_wait\_for\_pending\_transactions()** returns false. If *max\_delay* is set to zero, **pcie\_wait\_for\_pending\_transactions()** performs a single check; otherwise, this function may sleep while polling the transactions pending flag. **pcie\_wait\_for\_pending\_transactions** returns true if *dev* is not a PCI-express device.

## Device Configuration

The **pci\_enable\_busmaster()** function enables PCI bus mastering for the device *dev*, by setting the PCIM\_CMD\_BUSMASTEREN bit in the PCIR\_COMMAND register. The **pci\_disable\_busmaster()** function clears this bit.

The **pci\_enable\_io()** function enables memory or I/O port address decoding for the device *dev*, by setting the PCIM\_CMD\_MEMEN or PCIM\_CMD\_PORTEN bit in the PCIR\_COMMAND register appropriately. The **pci\_disable\_io()** function clears the appropriate bit. The *space* argument specifies which resource is affected; this can be either SYS\_RES\_MEMORY or SYS\_RES\_IOPORT as appropriate. Device drivers should generally not use these routines directly. The PCI bus will enable decoding automatically when a SYS\_RES\_MEMORY or SYS\_RES\_IOPORT resource is activated via **bus\_alloc\_resource(9)** or **bus\_activate\_resource(9)**.

The **pci\_get\_max\_payload()** function returns the current maximum TLP payload size in bytes for a PCI-express device. If the *dev* device is not a PCI-express device, **pci\_get\_max\_payload()** returns zero.

The **pci\_get\_max\_read\_req()** function returns the current maximum read request size in bytes for a PCI-express device. If the *dev* device is not a PCI-express device, **pci\_get\_max\_read\_req()** returns zero.

The **pci\_set\_max\_read\_req()** sets the PCI-express maximum read request size for *dev*. The requested *size* may be adjusted, and **pci\_set\_max\_read\_req()** returns the actual size set in bytes. If the *dev* device is not a PCI-express device, **pci\_set\_max\_read\_req()** returns zero.

The **pci\_get\_powerstate()** function returns the current power state of the device *dev*. If the device does not support power management capabilities, then the default state of PCI\_POWERSTATE\_D0 is returned. The following power states are defined by PCI:

PCI_POWERSTATE_D0	State in which device is on and running. It is receiving full power from the system and delivering full functionality to the user.
PCI_POWERSTATE_D1	Class-specific low-power state in which device context may or may not be lost. Buses in this state cannot do anything to the bus, to force devices to lose context.
PCI_POWERSTATE_D2	Class-specific low-power state in which device context may or may not be lost. Attains greater power savings than PCI_POWERSTATE_D1. Buses in this state can cause devices to lose some context. Devices <i>must</i> be prepared for the bus to be in this state or higher.
PCI_POWERSTATE_D3	State in which the device is off and not running. Device context is lost, and power from the device can be removed.
PCI_POWERSTATE_UNKNOWN	State of the device is unknown.

The **pci\_set\_powerstate()** function is used to transition the device *dev* to the PCI power state *state*. If the device does not support power management capabilities or it does not support the specific power state *state*, then the function will fail with EOPNOTSUPP.

The **pci\_iov\_attach()** function is used to advertise that the given device (and associated device driver) supports PCI Single-Root I/O Virtualization (SR-IOV). A driver that supports SR-IOV must implement the PCI\_IOV\_INIT(9), PCI\_IOV\_ADD\_VF(9) and PCI\_IOV\_UNINIT(9) methods. This function should be called during the DEVICE\_ATTACH(9) method. If this function returns an error, it is recommended that the device driver still successfully attaches, but runs with SR-IOV disabled. The *pf\_schema* and *vf\_schema* parameters are used to define what device-specific configuration parameters the device driver accepts when SR-IOV is enabled for the Physical Function (PF) and for individual



Virtual Functions (VFs) respectively. See `pci_iov_schema(9)` for details on how to construct the schema. If either the `pf_schema` or `vf_schema` is invalid or specifies parameter names that conflict with parameter names that are already in use, **`pci_iov_attach()`** will return an error and SR-IOV will not be available on the PF device. If a driver does not accept configuration parameters for either the PF device or the VF devices, the driver must pass an empty schema for that device. The SR-IOV infrastructure takes ownership of the `pf_schema` and `vf_schema` and is responsible for freeing them. The driver must never free the schemas itself.

The **`pci_iov_attach_name()`** function is a variant of **`pci_iov_attach()`** that allows the name of the associated character device in `/dev/iov` to be specified by *fnt*. The **`pci_iov_attach()`** function uses the name of *dev* as the device name.

The **`pci_iov_detach()`** function is used to advise the SR-IOV infrastructure that the driver for the given device is attempting to detach and that all SR-IOV resources for the device must be released. This function must be called during the `DEVICE_DETACH(9)` method if **`pci_iov_attach()`** was successfully called on the device and **`pci_iov_detach()`** has not subsequently been called on the device and returned no error. If this function returns an error, the `DEVICE_DETACH(9)` method must fail and return an error, as detaching the PF driver while VF devices are active would cause system instability. This function is safe to call and will always succeed if **`pci_iov_attach()`** previously failed with an error on the given device, or if **`pci_iov_attach()`** was never called on the device.

The **`pci_save_state()`** and **`pci_restore_state()`** functions can be used by a device driver to save and restore standard PCI config registers. The **`pci_save_state()`** function must be invoked while the device has valid state before **`pci_restore_state()`** can be used. If the device is not in the fully-powered state (`PCI_POWERSTATE_D0`) when **`pci_restore_state()`** is invoked, then the device will be transitioned to `PCI_POWERSTATE_D0` before any config registers are restored.

The **`pcie_flr()`** function requests a Function Level Reset (FLR) of *dev*. If *dev* is not a PCI-express device or does not support Function Level Resets via the PCI-express device control register, false is returned. Pending transactions are drained by disabling busmastering and calling **`pcie_wait_for_pending_transactions()`** before resetting the device. The *max\_delay* argument specifies the maximum timeout to wait for pending transactions as described for **`pcie_wait_for_pending_transactions()`**. If **`pcie_wait_for_pending_transactions()`** fails with a timeout and *force* is false, busmastering is re-enabled and false is returned. If **`pcie_wait_for_pending_transactions()`** fails with a timeout and *force* is true, the device is reset despite the timeout. After the reset has been requested, **`pcie_flr`** sleeps for at least 100 milliseconds before returning true. Note that **`pcie_flr`** does not save and restore any state around the reset. The caller should save and restore state as needed.

### Message Signaled Interrupts

Message Signaled Interrupts (MSI) and Enhanced Message Signaled Interrupts (MSI-X) are PCI

capabilities that provide an alternate method for PCI devices to signal interrupts. The legacy INTx interrupt is available to PCI devices as a `SYS_RES_IRQ` resource with a resource ID of zero. MSI and MSI-X interrupts are available to PCI devices as one or more `SYS_RES_IRQ` resources with resource IDs greater than zero. A driver must ask the PCI bus to allocate MSI or MSI-X interrupts using **`pci_alloc_msi()`** or **`pci_alloc_msix()`** before it can use MSI or MSI-X `SYS_RES_IRQ` resources. A driver is not allowed to use the legacy INTx `SYS_RES_IRQ` resource if MSI or MSI-X interrupts have been allocated, and attempts to allocate MSI or MSI-X interrupts will fail if the driver is currently using the legacy INTx `SYS_RES_IRQ` resource. A driver is only allowed to use either MSI or MSI-X, but not both.

The **`pci_msi_count()`** function returns the maximum number of MSI messages supported by the device *dev*. If the device does not support MSI, then **`pci_msi_count()`** returns zero.

The **`pci_alloc_msi()`** function attempts to allocate *\*count* MSI messages for the device *dev*. The **`pci_alloc_msi()`** function may allocate fewer messages than requested for various reasons including requests for more messages than the device *dev* supports, or if the system has a shortage of available MSI messages. On success, *\*count* is set to the number of messages allocated and **`pci_alloc_msi()`** returns zero. The `SYS_RES_IRQ` resources for the allocated messages will be available at consecutive resource IDs beginning with one. If **`pci_alloc_msi()`** is not able to allocate any messages, it returns an error. Note that MSI only supports message counts that are powers of two; requests to allocate a non-power of two count of messages will fail.

The **`pci_release_msi()`** function is used to release any allocated MSI or MSI-X messages back to the system. If any MSI or MSI-X `SYS_RES_IRQ` resources are allocated by the driver or have a configured interrupt handler, this function will fail with `EBUSY`. The **`pci_release_msi()`** function returns zero on success and an error on failure.

The **`pci_msix_count()`** function returns the maximum number of MSI-X messages supported by the device *dev*. If the device does not support MSI-X, then **`pci_msix_count()`** returns zero.

The **`pci_msix_pba_bar()`** function returns the offset in configuration space of the Base Address Register (BAR) containing the MSI-X Pending Bit Array (PBA) for device *dev*. The returned value can be used as the resource ID with `bus_alloc_resource(9)` and `bus_release_resource(9)` to allocate the BAR. If the device does not support MSI-X, then **`pci_msix_pba_bar()`** returns -1.

The **`pci_msix_table_bar()`** function returns the offset in configuration space of the BAR containing the MSI-X vector table for device *dev*. The returned value can be used as the resource ID with `bus_alloc_resource(9)` and `bus_release_resource(9)` to allocate the BAR. If the device does not support MSI-X, then **`pci_msix_table_bar()`** returns -1.

The **pci\_alloc\_msix()** function attempts to allocate *\*count* MSI-X messages for the device *dev*. The **pci\_alloc\_msix()** function may allocate fewer messages than requested for various reasons including requests for more messages than the device *dev* supports, or if the system has a shortage of available MSI-X messages. On success, *\*count* is set to the number of messages allocated and **pci\_alloc\_msix()** returns zero. For MSI-X messages, the resource ID for each SYS\_RES\_IRQ resource identifies the index in the MSI-X table of the corresponding message. A resource ID of one maps to the first index of the MSI-X table; a resource ID two identifies the second index in the table, etc. The **pci\_alloc\_msix()** function assigns the *\*count* messages allocated to the first *\*count* table indices. If **pci\_alloc\_msix()** is not able to allocate any messages, it returns an error. Unlike MSI, MSI-X does not require message counts that are powers of two.

The BARs containing the MSI-X vector table and PBA must be allocated via **bus\_alloc\_resource(9)** before calling **pci\_alloc\_msix()** and must not be released until after calling **pci\_release\_msi()**. Note that the vector table and PBA may be stored in the same BAR or in different BARs.

The **pci\_pending\_msix()** function examines the *dev* device's PBA to determine the pending status of the MSI-X message at table index *index*. If the indicated message is pending, this function returns a non-zero value; otherwise, it returns zero. Passing an invalid *index* to this function will result in undefined behavior.

As mentioned in the description of **pci\_alloc\_msix()**, MSI-X messages are initially assigned to the first *N* table entries. A driver may use a different distribution of available messages to table entries via the **pci\_remap\_msix()** function. Note that this function must be called after a successful call to **pci\_alloc\_msix()** but before any of the SYS\_RES\_IRQ resources are allocated. The **pci\_remap\_msix()** function returns zero on success, or an error on failure.

The *vectors* array should contain *count* message vectors. The array maps directly to the MSI-X table in that the first entry in the array specifies the message used for the first entry in the MSI-X table, the second entry in the array corresponds to the second entry in the MSI-X table, etc. The vector value in each array index can either be zero to indicate that no message should be assigned to the corresponding MSI-X table entry, or it can be a number from one to *N* (where *N* is the count returned from the previous call to **pci\_alloc\_msix()**) to indicate which of the allocated messages should be assigned to the corresponding MSI-X table entry.

If **pci\_remap\_msix()** succeeds, each MSI-X table entry with a non-zero vector will have an associated SYS\_RES\_IRQ resource whose resource ID corresponds to the table index as described above for **pci\_alloc\_msix()**. MSI-X table entries that with a vector of zero will not have an associated SYS\_RES\_IRQ resource. Additionally, if any of the original messages allocated by **pci\_alloc\_msix()** are not used in the new distribution of messages in the MSI-X table, they will be released automatically. Note that if a driver wishes to use fewer messages than were allocated by **pci\_alloc\_msix()**, the driver

must use a single, contiguous range of messages beginning with one in the new distribution. The **pci\_remap\_msix()** function will fail if this condition is not met.

### Device Events

The *pci\_add\_device* event handler is invoked every time a new PCI device is added to the system. This includes the creation of Virtual Functions via SR-IOV.

The *pci\_delete\_device* event handler is invoked every time a PCI device is removed from the system.

Both event handlers pass the *device\_t* object of the relevant PCI device as *dev* to each callback function. Both event handlers are invoked while *dev* is unattached but with valid instance variables.

### SEE ALSO

pci(4), pciconf(8), bus\_alloc\_resource(9), bus\_dma(9), bus\_release\_resource(9), bus\_setup\_intr(9), bus\_teardown\_intr(9), devclass(9), device(9), driver(9), eventhandler(9), rman(9)

*NewBus, FreeBSD Developers' Handbook*, <https://docs.freebsd.org/en/books/developers-handbook/>.

Shanley and Anderson, *PCI System Architecture*, Addison-Wesley, 2nd Edition, ISBN 0-201-30974-2.

### AUTHORS

This manual page was written by Bruce M Simpson <[bms@FreeBSD.org](mailto:bms@FreeBSD.org)> and John Baldwin <[jhb@FreeBSD.org](mailto:jhb@FreeBSD.org)>.

### BUGS

The kernel PCI code has a number of references to "slot numbers". These do not refer to the geographic location of PCI devices, but to the device number assigned by the combination of the PCI IDSEL mechanism and the platform firmware. This should be taken note of when working with the kernel PCI code.

The PCI bus driver should allocate the MSI-X vector table and PBA internally as necessary rather than requiring the caller to do so.