## NAME

aio - asynchronous I/O

# DESCRIPTION

The **aio** facility provides system calls for asynchronous I/O. Asynchronous I/O operations are not completed synchronously by the calling thread. Instead, the calling thread invokes one system call to request an asynchronous I/O operation. The status of a completed request is retrieved later via a separate system call.

Asynchronous I/O operations on some file descriptor types may block an AIO daemon indefinitely resulting in process and/or system hangs. Operations on these file descriptor types are considered "unsafe" and disabled by default. They can be enabled by setting the *vfs.aio.enable\_unsafe* sysctl node to a non-zero value.

Asynchronous I/O operations on sockets, raw disk devices, and regular files on local filesystems do not block indefinitely and are always enabled.

The **aio** facility uses kernel processes (also known as AIO daemons) to service most asynchronous I/O requests. These processes are grouped into pools containing a variable number of processes. Each pool will add or remove processes to the pool based on load. Pools can be configured by sysctl nodes that define the minimum and maximum number of processes as well as the amount of time an idle process will wait before exiting.

One pool of AIO daemons is used to service asynchronous I/O requests for sockets. These processes are named "soaiod<N>". The following syscel nodes are used with this pool:

### kern.ipc.aio.num\_procs

The current number of processes in the pool.

### kern.ipc.aio.target\_procs

The minimum number of processes that should be present in the pool.

### kern.ipc.aio.max\_procs

The maximum number of processes permitted in the pool.

### kern.ipc.aio.lifetime

The amount of time a process is permitted to idle in clock ticks. If a process is idle for this amount of time and there are more processes in the pool than the target minimum, the process will exit.

A second pool of AIO daemons is used to service all other asynchronous I/O requests except for I/O requests to raw disks. These processes are named "aiod<N>". The following sysctl nodes are used with this pool:

## vfs.aio.num\_aio\_procs

The current number of processes in the pool.

## vfs.aio.target\_aio\_procs

The minimum number of processes that should be present in the pool.

## vfs.aio.max\_aio\_procs

The maximum number of processes permitted in the pool.

## vfs.aio.aiod\_lifetime

The amount of time a process is permitted to idle in clock ticks. If a process is idle for this amount of time and there are more processes in the pool than the target minimum, the process will exit.

Asynchronous I/O requests for raw disks are queued directly to the disk device layer after temporarily wiring the user pages associated with the request. These requests are not serviced by any of the AIO daemon pools.

Several limits on the number of asynchronous I/O requests are imposed both system-wide and perprocess. These limits are configured via the following syscels:

# vfs.aio.max\_buf\_aio

The maximum number of queued asynchronous I/O requests for raw disks permitted for a single process. Asynchronous I/O requests that have completed but whose status has not been retrieved via aio\_return(2) or aio\_waitcomplete(2) are not counted against this limit.

# vfs.aio.num\_buf\_aio

The number of queued asynchronous I/O requests for raw disks system-wide.

### vfs.aio.max\_aio\_queue\_per\_proc

The maximum number of asynchronous I/O requests for a single process serviced concurrently by the default AIO daemon pool.

### vfs.aio.max\_aio\_per\_proc

The maximum number of outstanding asynchronous I/O requests permitted for a single process. This includes requests that have not been serviced, requests currently being serviced, and

requests that have completed but whose status has not been retrieved via aio\_return(2) or aio\_waitcomplete(2).

## vfs.aio.num\_queue\_count

The number of outstanding asynchronous I/O requests system-wide.

## vfs.aio.max\_aio\_queue

The maximum number of outstanding asynchronous I/O requests permitted system-wide.

Asynchronous I/O control buffers should be zeroed before initializing individual fields. This ensures all fields are initialized.

All asynchronous I/O control buffers contain a *sigevent* structure in the *aio\_sigevent* field which can be used to request notification when an operation completes.

For SIGEV\_KEVENT notifications, the *sigevent*'s *sigev\_notify\_kqueue* field should contain the descriptor of the kqueue that the event should be attached to, its *sigev\_notify\_kevent\_flags* field may contain EV\_ONESHOT, EV\_CLEAR, and/or EV\_DISPATCH, and its *sigev\_notify* field should be set to SIGEV\_KEVENT. The posted kevent will contain:

Member	Value
ident	asynchronous I/O control buffer pointer
filter	EVFILT_AIO
flags	EV_EOF
udata	value stored in <i>aio_sigevent.sigev_value</i>

For SIGEV\_SIGNO and SIGEV\_THREAD\_ID notifications, the information for the queued signal will include SI\_ASYNCIO in the *si\_code* field and the value stored in *sigevent.sigev\_value* in the *si\_value* field.

For SIGEV\_THREAD notifications, the value stored in *aio\_sigevent.sigev\_value* is passed to the *aio\_sigevent.sigev\_notify\_function* as described in sigevent(3).

# SEE ALSO

aio\_cancel(2), aio\_error(2), aio\_read(2), aio\_readv(2), aio\_return(2), aio\_suspend(2), aio\_waitcomplete(2), aio\_write(2), aio\_writev(2), lio\_listio(2), sigevent(3), sysctl(8)

### HISTORY

The **aio** facility appeared as a kernel option in FreeBSD 3.0. The **aio** kernel module appeared in FreeBSD 5.0. The **aio** facility was integrated into all kernels in FreeBSD 11.0.