

**NAME**

**sched\_4bsd** - 4.4BSD scheduler

**SYNOPSIS**

**options SCHED\_4BSD**

**DESCRIPTION**

The **sched\_4bsd** scheduler is the traditional system scheduler, providing both high throughput and solid interactive response in the presence of load.

The following sysctls are relevant to the operation of **sched\_4bsd**:

*kern.sched.name*

This read-only sysctl reports the name of the active scheduler.

*kern.sched.quantum*

This read-write sysctl reports or sets the length of the quantum (in micro-seconds) granted to a thread.

*kern.sched.ipiwakeupt.enabled*

This read-write sysctl sets whether or not the scheduler will generate an inter-processor interrupt (IPI) to an idle CPU when a thread is woken up. Otherwise, idle CPUs will wait until the next clock tick before looking for new work.

*kern.sched.preemption*

This read-only sysctl reports whether or not the kernel is configured to support preemption, which reduces the latency to run lower priority threads on wakeup.

Some sysctls will be available only on systems supporting SMP.

**SEE ALSO**

sched\_ule(4), sysctl(8)

**HISTORY**

The **sched\_4bsd** scheduler has been present, in various forms, since the inception of BSD.

**BUGS**

While a highly robust and time-tested scheduler, **sched\_4bsd** lacks specific knowledge of how to schedule advantageously in non-symmetric processor configurations, such as hyper-threading.