

NAME

SMP - description of the FreeBSD Symmetric Multi-Processor kernel

SYNOPSIS

options SMP

DESCRIPTION

The **SMP** kernel implements symmetric multi-processor support.

SMP support can be disabled by setting the loader tunable *kern.smp.disabled* to 1.

The number of CPUs detected by the system is available in the read-only sysctl variable *hw.ncpu*.

The number of online threads per CPU core is available in the read-only sysctl variable *kern.smp.threads_per_core*. The number of physical CPU cores detected by the system is available in the read-only sysctl variable *kern.smp.cores*.

FreeBSD allows specific CPUs on a multi-processor system to be disabled. This can be done using the *hint.lapic.X.disabled* tunable, where X is the APIC ID of a CPU. Setting this tunable to 1 will result in the corresponding CPU being disabled.

FreeBSD supports simultaneous multithreading on x86 and powerpc platforms. On x86, the logical CPUs can be disabled by setting the *machdep.hyperthreading_allowed* tunable to zero.

The sched_ule(4) scheduler implements CPU topology detection and adjusts the scheduling algorithms to make better use of modern multi-core CPUs. The sysctl variable *kern.sched.topology_spec* reflects the detected CPU hardware in a parsable XML format. The top level XML tag is <groups>, which encloses one or more <group> tags containing data about individual CPU groups. A CPU group contains CPUs that are detected to be "close" together, usually by being cores in a single multi-core processor. Attributes available in a <group> tag are "level", corresponding to the nesting level of the CPU group and "cache-level", corresponding to the level of CPU caches shared by the CPUs in the group. The <group> tag contains the <cpu> and <flags> tags. The <cpu> tag describes CPUs in the group. Its attributes are "count", corresponding to the number of CPUs in the group and "mask", corresponding to the integer binary mask in which each bit position set to 1 signifies a CPU belonging to the group. The contents (CDATA) of the <cpu> tag is the comma-delimited list of CPU indexes (derived from the "mask" attribute). The <flags> tag contains special tags (if any) describing the relation of the CPUs in the group. The possible flags are currently "HTT" and "SMT", corresponding to the various implementations of hardware multithreading. An example topology_spec output for a system consisting of two quad-core processors is:

```

<groups>
  <group level="1" cache-level="0">
    <cpu count="8" mask="0xff">0, 1, 2, 3, 4, 5, 6, 7</cpu>
    <flags></flags>
    <children>
      <group level="2" cache-level="0">
        <cpu count="4" mask="0xf">0, 1, 2, 3</cpu>
        <flags></flags>
      </group>
      <group level="2" cache-level="0">
        <cpu count="4" mask="0xf0">4, 5, 6, 7</cpu>
        <flags></flags>
      </group>
    </children>
  </group>
</groups>

```

This information is used internally by the kernel to schedule related tasks on CPUs that are closely grouped together.

COMPATIBILITY

Support for multi-processor systems is present for all Tier-1 and Tier-2 architectures on FreeBSD. Currently, this includes x86, powerpc, mips, arm and arm64. Support is enabled using **options SMP**. It is permissible to use the SMP kernel configuration on non-SMP hardware.

I386 NOTES

For i386 systems, the **SMP** kernel supports motherboards that follow the Intel MP specification, version 1.4. In addition to **options SMP**, i386 also requires **device apic**. The `mptable(1)` command may be used to view the status of multi-processor support.

SEE ALSO

`cpuset(1)`, `mptable(1)`, `sched_4bsd(4)`, `sched_ule(4)`, `loader(8)`, `sysctl(8)`, `condvar(9)`, `msleep(9)`, `mtx_pool(9)`, `mutex(9)`, `rwlock(9)`, `sema(9)`, `sx(9)`

HISTORY

The **SMP** kernel's early history is not (properly) recorded. It was developed in a separate CVS branch until April 26, 1997, at which point it was merged into 3.0-current. By this date 3.0-current had already been merged with Lite2 kernel code.

FreeBSD 5.0 introduced support for a host of new synchronization primitives, and a move towards fine-

grained kernel locking rather than reliance on a Giant kernel lock. The SMPng Project relied heavily on the support of BSDi, who provided reference source code from the fine-grained SMP implementation found in BSD/OS.

FreeBSD 5.0 also introduced support for SMP on the sparc64 architecture.

AUTHORS

Steve Passe <fsmp@FreeBSD.org>

CAVEATS

The *kern.smp.threads_per_core* and *kern.smp.cores* sysctl variables are provided as a best-effort guess. If an architecture or platform adds SMT and FreeBSD has not yet implemented detection, the reported values may be inaccurate. In this case, *kern.smp.threads_per_core* will report 1 and *kern.smp.cores* will report the same value as *hw.ncpu*.