

# [PATCH][RFC] Add default CPU topology information

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Kernel/2008-03/msg11959.html>

---

- *From:* Ben Hutchings <[bhutchings@xxxxxxxxxxxxxxxx](mailto:bhutchings@xxxxxxxxxxxxxxxx)>
  - *Date:* Mon, 31 Mar 2008 12:44:13 +0100
- 

Not all architectures and configurations define CPU topology information. This can result in an empty topology directory in sysfs, and requires in-kernel users to protect all uses with #ifdef – see <http://marc.info/?l=linux-netdev&m=120639033904472&w=2>.

The documentation of CPU topology specifies what the defaults should be if only partial information is available from the hardware. So we can provide these defaults as a fallback.

This patch:

- Adds default definitions of the 4 topology macros to `include/asm-generic/topology.h`.
- Changes `include/asm-*/topology.h` to include `<asm-generic/topology.h>` unconditionally.
- Changes `drivers/base/topology.c` to use the topology macros unconditionally and to cope with definitions that aren't lvalues.
- Updates documentation accordingly.

Signed-off-by: Ben Hutchings <[bhutchings@xxxxxxxxxxxxxxxx](mailto:bhutchings@xxxxxxxxxxxxxxxx)>  
diff --git a/Documentation/cputopology.txt b/Documentation/cputopology.txt  
index b61cb95..c035481 100644  
--- a/Documentation/cputopology.txt  
+++ b/Documentation/cputopology.txt  
@@ -14,9+14,8 @@ represent the thread siblings to cpu X in the same physical package;  
To implement it in an architecture-neutral way, a new source file, `drivers/base/topology.c`, is to export the 4 attributes.

- If one architecture wants to support this feature, it just needs to
- implement 4 defines, typically in file `include/asm-XXX/topology.h`.
- The 4 defines are:  
+If one architecture wants to support this feature, it must define  
+some of these macros in `include/asm-XXX/topology.h`:  
#define topology\_physical\_package\_id(cpu)  
#define topology\_core\_id(cpu)  
#define topology\_thread\_siblings(cpu)  
@@ -25,17+24,10 @@ The 4 defines are:  
The type of `**_id` is int.

## [PATCH][RFC] Add default CPU topology information

The type of siblings is cpumask\_t.

- To be consistent on all architectures, the 4 attributes should have
- default values if their values are unavailable. Below is the rule.
- 1) physical\_package\_id: If cpu has no physical package id, -1 is the
- default value.
- 2) core\_id: If cpu doesn't support multi-core, its core id is 0.
- 3) thread\_siblings: Just include itself, if the cpu doesn't support
- HT/multi-thread.
- 4) core\_siblings: Just include itself, if the cpu doesn't support
- multi-core and HT/Multi-thread.

-

-So be careful when declaring the 4 defines in include/asm-XXX/topology.h.

-

-If an attribute isn't defined on an architecture, it won't be exported.

-

+To be consistent on all architectures, include/asm-generic/topology.h

+provides default definitions for any of the above macros that are

+not already defined:

+1) physical\_package\_id: -1

+2) core\_id: 0

+3) thread\_siblings: just the given CPU

+4) core\_siblings: just the given CPU

diff --git a/drivers/base/topology.c b/drivers/base/topology.c

index e1d3ad4..dfcb2e2 100644

--- a/drivers/base/topology.c

+++ b/drivers/base/topology.c

@@ -45,47 +45,28 @@ static ssize\_t show\_###name(struct sys\_device \*dev, char \*buf) \

{ \

ssize\_t len = -1; \

unsigned int cpu = dev->id; \

- len = cpumask\_scnprintf(buf, NR\_CPUS+1, topology\_###name(cpu)); \

+ cpumask\_t siblings = topology\_###name(cpu); \

+ len = cpumask\_scnprintf(buf, NR\_CPUS+1, siblings); \

return (len + sprintf(buf + len, "\n")); \

}

-#ifdef topology\_physical\_package\_id

define\_id\_show\_func(physical\_package\_id);

define\_one\_ro(physical\_package\_id);

-#define ref\_physical\_package\_id\_attr &attr\_physical\_package\_id.attr,

-#else

-#define ref\_physical\_package\_id\_attr

-#endif

-#ifdef topology\_core\_id

define\_id\_show\_func(core\_id);

define\_one\_ro(core\_id);

-#define ref\_core\_id\_attr &attr\_core\_id.attr,

-#else

-#define ref\_core\_id\_attr

[PATCH][RFC] Add default CPU topology information

[PATCH][RFC] Add default CPU topology information

```

-#endif

-#ifdef topology_thread_siblings
define_siblings_show_func(thread_siblings);
define_one_ro(thread_siblings);
-#define ref_thread_siblings_attr &attr_thread_siblings.attr,
-#else
-#define ref_thread_siblings_attr
-#endif

-#ifdef topology_core_siblings
define_siblings_show_func(core_siblings);
define_one_ro(core_siblings);
-#define ref_core_siblings_attr &attr_core_siblings.attr,
-#else
-#define ref_core_siblings_attr
-#endif

static struct attribute *default_attrs[] = {
- ref_physical_package_id_attr
- ref_core_id_attr
- ref_thread_siblings_attr
- ref_core_siblings_attr
+ &attr_physical_package_id.attr,
+ &attr_core_id.attr,
+ &attr_thread_siblings.attr,
+ &attr_core_siblings.attr,
NULL
};

diff --git a/include/asm-alpha/topology.h b/include/asm-alpha/topology.h
index 420ccde..e59ec8a 100644
--- a/include/asm-alpha/topology.h
+++ b/include/asm-alpha/topology.h
@@ -41,8 +41,8 @@ static inline cpumask_t node_to_cpumask(int node)

#define pcibus_to_cpumask(bus) (cpu_online_map)

-#else /* CONFIG_NUMA */
-# include <asm-generic/topology.h>
-#endif /* !CONFIG_NUMA */
+#endif /* CONFIG_NUMA */
+
+#include <asm-generic/topology.h>

#endif /* _ASM_ALPHA_TOPOLOGY_H */
diff --git a/include/asm-frv/topology.h b/include/asm-frv/topology.h
index abe7298..9427243 100644
--- a/include/asm-frv/topology.h
+++ b/include/asm-frv/topology.h
@@ -5,10 +5,8 @@

```

[PATCH][RFC] Add default CPU topology information

```
#error NUMA not supported yet

-#else /* !CONFIG_NUMA */
+#endif /* CONFIG_NUMA */

#include <asm-generic/topology.h>

-#endif /* CONFIG_NUMA */
-
#endif /* _ASM_TOPOLOGY_H */
diff --git a/include/asm-generic/topology.h b/include/asm-generic/topology.h
index 342a2a0..373502b 100644
--- a/include/asm-generic/topology.h
+++ b/include/asm-generic/topology.h
@@ -27,6 +27,8 @@
#ifndef _ASM_GENERIC_TOPOLOGY_H
#define _ASM_GENERIC_TOPOLOGY_H

+#include <linux/cpumask.h>
+
/* Other architectures wishing to use this simple topology API should fill
in the below functions as appropriate in their own <asm/topology.h> file. */
#ifndef cpu_to_node
@@ -52,4 +54,17 @@
)
#endif

+#ifndef topology_physical_package_id
+#define topology_physical_package_id(cpu) ((void)(cpu), -1)
+#endif
+#ifndef topology_core_id
+#define topology_core_id(cpu) ((void)(cpu), 0)
+#endif
+#ifndef topology_thread_siblings
+#define topology_thread_siblings(cpu) cpumask_of_cpu(cpu)
+#endif
+#ifndef topology_core_siblings
+#define topology_core_siblings(cpu) cpumask_of_cpu(cpu)
+#endif
+
#endif /* _ASM_GENERIC_TOPOLOGY_H */
diff --git a/include/asm-mips/topology.h b/include/asm-mips/topology.h
index 259145e..ce837a4 100644
--- a/include/asm-mips/topology.h
+++ b/include/asm-mips/topology.h
@@ -8,10 +8,10 @@
#ifndef __ASM_TOPOLOGY_H
#define __ASM_TOPOLOGY_H

-#include <topology.h>
```

[PATCH][RFC] Add default CPU topology information

[PATCH][RFC] Add default CPU topology information

```
-
#ifdef CONFIG_SMP
#define smt_capable() (smp_num_siblings > 1)
#endif

+#include <asm-generic/topology.h>
+
#endif /* __ASM_TOPOLOGY_H */
diff --git a/include/asm-powerpc/topology.h b/include/asm-powerpc/topology.h
index ca23b68..e587788 100644
--- a/include/asm-powerpc/topology.h
+++ b/include/asm-powerpc/topology.h
@@ -96,9 +96,6 @@ static inline void sysfs_remove_device_from_node(struct sys_device *dev,
{
}

-
-#include <asm-generic/topology.h>
-
#endif /* CONFIG_NUMA */

#ifdef CONFIG_SMP
@@ -112,5 +109,7 @@ static inline void sysfs_remove_device_from_node(struct sys_device *dev,
#endif
#endif

+#include <asm-generic/topology.h>
+
#endif /* __KERNEL__ */
#endif /* _ASM_POWERPC_TOPOLOGY_H */
diff --git a/include/asm-x86/topology.h b/include/asm-x86/topology.h
index 8af05a9..7f31b8d 100644
--- a/include/asm-x86/topology.h
+++ b/include/asm-x86/topology.h
@@ -165,11 +165,7 @@ extern int __node_distance(int, int);
#define node_distance(a, b) __node_distance(a, b)
#endif

-#else /* CONFIG_NUMA */
-
-#include <asm-generic/topology.h>
-
-#endif
+#endif /* CONFIG_NUMA */

extern cpumask_t cpu_coregroup_map(int cpu);

@@ -185,4 +181,6 @@ extern cpumask_t cpu_coregroup_map(int cpu);
#define smt_capable() (smp_num_siblings > 1)
#endif
```

[PATCH][RFC] Add default CPU topology information

```
+#include <asm-generic/topology.h>  
+  
#endif
```

--

Ben Hutchings, Senior Software Engineer, Solarflare Communications  
Not speaking for my employer; that's the marketing department's job.

--

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in  
the body of a message to majordomo@xxxxxxxxxxxxxxxxx

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>