

Re: [Patch:001/004]Unify pxm_to_node id ver.3.(generic code)

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-03/msg09707.html>

- *From:* Andrew Morton <akpm@xxxxxxxx>
 - *Date:* Tue, 28 Mar 2006 13:07:36 -0800
-

Yasunori Goto <y-goto@xxxxxxxxxxxxxxxx> wrote:

```
+/* Proximity bitmap length */
+#ifdef CONFIG_NR_NODES_CHANGABLE
+#define MAX_PXM_DOMAINS CONFIG_NR_NODES
+#else
+#define MAX_PXM_DOMAINS (256)
+#endif
```

I don't think we need CONFIG_NR_NODES_CHANGABLE (it is spelled "changeable", btw).

If the architecture wants to support changing of CONFIG_NR_NODES then it can permit CONFIG_NR_NODES to be changed in its Kconfig implementation.

If the architecture doesn't want to permit changing of CONFIG_NR_NODES then it should simply hardwire CONFIG_NR_NODES to the chosen value in its Kconfig.

So all architectures which use acpi_numa must implement CONFIG_NR_NODES.

In fact, it would probably make sense to require that all NUMA-supporting architectures implement CONFIG_NR_NODES.

Also, we already have NODES_SHIFT defined in include/asm-*/numnodes.h. What's the relationship between that and CONFIG_NR_NODES? It seems that we want to derive NODES_SHIFT from CONFIG_NR_NODES.

Was ia64's CONFIG_IA64_NR_NODES the best choice? Should ia64 instead have made NODES_SHIFT Kconfigurable, and derived its max_nr_nodes from that?

It's all a bit of a pickle.

I guess for now a suitable approach would be to make all numa-using architectures define CONFIG_NR_NODES, and to leave that rather

Re: [Patch:001/004]Unify pxm_to_node id ver.3.(generic code)

unpleasant-looking code in include/asm-ia64/numnodes.h as it is.

—

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

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

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