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

Source: http://linux.derkeiler.com/Mailing−Lists/Kernel/2006−03/msg09536.html

- *From*: Yasunori Goto <y−goto@xxxxxxxxxxxxxx>
- *Date*: Tue, 28 Mar 2006 19:16:56 +0900

This is new generic code for pxm_to_node_map and CONFIG_NR_NODES.

Signed−off−by: Yasunori Goto <y−goto@xxxxxxxxxxxxxx>

drivers/acpi/numa.c | 48 ++
include/acpi/acpi_numa.h | 23 ++++++++++++++++++++++
include/linux/acpi.h | 1
mm/Kconfig | 12 ++++++++++
4 files changed, 84 insertions(+)

Index: pxm_ver3/drivers/acpi/numa.c
===
−−− pxm_ver3.orig/drivers/acpi/numa.c 2006−03−28 14:10:02.867761158 +0900
+++ pxm_ver3/drivers/acpi/numa.c 2006−03−28 14:13:30.926352359 +0900
@@ −36,12 +36,60 @@
#define _COMPONENT ACPI_NUMA
ACPI_MODULE_NAME("numa")

+static nodemask_t nodes_found_map = NODE_MASK_NONE;
+#define PXM_INVAL −1
+#define NID_INVAL −1
+
+/* maps to convert between proximity domain and logical node ID */
+int __cpuinitdata pxm_to_node_map[MAX_PXM_DOMAINS]
+ = { [0 ... MAX_PXM_DOMAINS − 1] = NID_INVAL };
+int __cpuinitdata node_to_pxm_map[MAX_NUMNODES]
+ = { [0 ... MAX_NUMNODES − 1] = PXM_INVAL };
+
extern int __init acpi_table_parse_madt_family(enum acpi_table_id id,
unsigned long madt_size,
int entry_id,
acpi_madt_entry_handler handler,
unsigned int max_entries);

+int __cpuinit pxm_to_node(int pxm)
+{
+ if (pxm < 0)
+ return NID_INVAL;

```
+ return pxm_to_node_map[pxm];
+}
+
+int __cpuinit node_to_pxm(int node)
+{
+ if (node < 0)
+ return PXM_INVAL;
+ return node_to_pxm_map[node];
+}
+
+int __cpuinit acpi_map_pxm_to_node(int pxm)
+{
+ int node = pxm_to_node_map[pxm];
+
+ if (node < 0){
+ if (nodes_weight(nodes_found_map) >= MAX_NUMNODES)
+ return NID_INVAL;
+ node = first_unset_node(nodes_found_map);
+ pxm_to_node_map[pxm] = node;
+ node_to_pxm_map[node] = pxm;
+ node_set(node, nodes_found_map);
+ }
+
+ return node;
+}
+
+void __cpuinit acpi_unmap_pxm_to_node(int node)
+{
+ int pxm = node_to_pxm_map[node];
+ pxm_to_node_map[pxm] = NID_INVAL;
+ node_to_pxm_map[node] = PXM_INVAL;
+ node_clear(node, nodes_found_map);
+}
+
void __init acpi_table_print_srat_entry(acpi_table_entry_header * header)
{

Index: pxm_ver3/include/acpi/acpi_numa.h
===================================================================
--- /dev/null 1970-01-01 00:00:00.000000000 +0000
+++ pxm_ver3/include/acpi/acpi_numa.h 2006-03-28 14:13:30.927328921 +0900
@@ -0,0 +1,23 @@
+#ifndef __ACPI_NUMA_H
+#define __ACPI_NUMA_H
+
+#ifdef CONFIG_ACPI_NUMA
+#include <linux/kernel.h>
+
+/* Proximity bitmap length */
+#ifdef CONFIG_NR_NODES_CHANGABLE
+#define MAX_PXM_DOMAINS CONFIG_NR_NODES
```

```
+#else
+#define MAX_PXM_DOMAINS (256)
+#endif
+
+extern int __cpuinitdata pxm_to_node_map[MAX_PXM_DOMAINS];
+extern int __cpuinitdata node_to_pxm_map[MAX_NUMNODES];
+
+extern int __cpuinit pxm_to_node(int);
+extern int __cpuinit node_to_pxm(int);
+extern int __cpuinit acpi_map_pxm_to_node(int);
+extern void __cpuinit acpi_unmap_pxm_to_node(int);
+
+#endif /* CONFIG_ACPI_NUMA */
+#endif /* __ACP_NUMA_H */
Index: pxm_ver3/include/linux/acpi.h
===================================================================
--- pxm_ver3.orig/include/linux/acpi.h 2006-03-28 14:10:02.867761158 +0900
+++ pxm_ver3/include/linux/acpi.h 2006-03-28 14:24:38.740797303 +0900
@@ -38,6 +38,7 @@
#include <acpi/acpi.h>
#include <acpi/acpi_bus.h>
#include <acpi/acpi_drivers.h>
+#include <acpi/acpi_numa.h>
#include <asm/acpi.h>


Index: pxm_ver3/mm/Kconfig
===================================================================
--- pxm_ver3.orig/mm/Kconfig 2006-03-28 14:24:38.009352000 +0900
+++ pxm_ver3/mm/Kconfig 2006-03-28 14:24:53.320875250 +0900
@@ -91,6 +91,18 @@ config HAVE_MEMORY_PRESENT
depends on ARCH_HAVE_MEMORY_PRESENT || SPARSEMEM

#
+# NR_NODES is to configure NODES_SHIFT
+#
+config NR_NODES
+ int "Maximum number of NODEs (256-1024)"
+ range 256 1024
+ depends on NEED_MULTIPLE_NODES && NR_NODES_CHANGABLE
+ default "256"
+ help
+ This option specifies the maximum number of nodes in your SSI system.
+ If in doubt, use the default.
+
+#
# SPARSEMEM_EXTREME (which is the default) does some bootmem
# allocations when memory_present() is called. If this can not
# be done on your architecture, select this option. However,

--
```

Yasunori Goto

–