

Re: [Lse-tech] [PATCH] new bitmap list format (for cpusets)

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-08/2183.html>

From: Paul Jackson (pj_at_sgi.com)

Date: 08/09/04

Date: Mon, 9 Aug 2004 01:01:06 -0700
To: "Martin J. Bligh" <mbligh@aracnet.com>

I was looking at this bitmap list format patch over the weekend, and came to the conclusion that the basic list format, as in the example:

0,3,5,8-15

was a valuable improvement over a fixed length hex mask, but that on the other hand the support for:

the prefix characters '=', '-', '+', or '!'

was fluff, that few would learn to use, and fewer find essential.

So I redid the bitmap list format patch, removing the prefix character support, and making another pass at compacting the input 'write' side code in `bitmap_parselist()`.

The kernel text costs in bytes for these two patches, on an i386 build, are now:

```
bitmap lists: 592
cpusets: 7718
-----
total: 8310
```

Here's the new bitmap list patch. It applies to 2.6.8-rc2-mm2. It replaces the earlier bitmap list patch, that began this thread on August 5, 2004.

=====

A bitmap print and parse format that provides lists of ranges of numbers, to be first used for by cpusets (next patch).

Cpusets provide a way to manage subsets of CPUs and Memory Nodes for scheduling and memory placement, via a new virtual file system,

Linux-Kernel: Re: [Lse-tech] [PATCH] new bitmap list format (for cpusets)

usually mounted at /dev/cpuset. Manipulation of cpusets can be done directly via this file system, from the shell.

However, manipulating 512 bit cpumasks or 256 bit nodemasks (which will get bigger) via hex mask strings is painful for humans.

The intention is to provide a format for the cpu and memory mask files in /dev/cpusets that will stand the test of time. This format is supported by a couple of new lib/bitmap.c routines, for printing and parsing these strings. Wrappers for cpumask and nodemask are provided.

```
include/linux/bitmap.h | 8 +++
include/linux/cpumask.h | 22 ++++++++
include/linux/nodemask.h | 22 ++++++++
lib/bitmap.c | 103 +++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++
4 files changed, 150 insertions(+), 5 deletions(-)
```

Signed-off-by: Paul Jackson <pj@sgi.com>

Index: 2.6.8-rc2-mm2/include/linux/bitmap.h

```
----- 2.6.8-rc2-mm2.orig/include/linux/bitmap.h 2004-08-08 23:17:35.000000000 -0700
+++ 2.6.8-rc2-mm2/include/linux/bitmap.h 2004-08-08 23:24:57.000000000 -0700
@@ -41,7 +41,9 @@
 * bitmap_shift_right(dst, src, n, nbits) *dst = *src >> n
 * bitmap_shift_left(dst, src, n, nbits) *dst = *src << n
 * bitmap_scnprintf(buf, len, src, nbits) Print bitmap src to buf
- * bitmap_parse(ubuf, ulen, dst, nbits) Parse bitmap dst from buf
+ * bitmap_parse(ubuf, ulen, dst, nbits) Parse bitmap dst from user buf
+ * bitmap_scnlistprintf(buf, len, src, nbits) Print bitmap src as list to buf
+ * bitmap_parselist(buf, dst, nbits) Parse bitmap dst from list
 */

/*
@@ -98,6 +100,10 @@ extern int bitmap_scnprintf(char *buf, u
const unsigned long *src, int nbits);
extern int bitmap_parse(const char __user *ubuf, unsigned int ulen,
unsigned long *dst, int nbits);
+extern int bitmap_scnlistprintf(char *buf, unsigned int len,
+ const unsigned long *src, int nbits);
+extern int bitmap_parselist(const char *buf, unsigned long *maskp,
+ int nmaskbits);
extern int bitmap_find_free_region(unsigned long *bitmap, int bits, int order);
extern void bitmap_release_region(unsigned long *bitmap, int pos, int order);
extern int bitmap_allocate_region(unsigned long *bitmap, int pos, int order);
```

Index: 2.6.8-rc2-mm2/include/linux/cpumask.h

```
----- 2.6.8-rc2-mm2.orig/include/linux/cpumask.h 2004-08-08 23:17:35.000000000 -0700
+++ 2.6.8-rc2-mm2/include/linux/cpumask.h 2004-08-08 23:24:57.000000000 -0700
@@ -10,6 +10,8 @@
*
```

Re: [Lse-tech] [PATCH] new bitmap list format (for cpusets)

Linux-Kernel: Re: [Lse-tech] [PATCH] new bitmap list format (for cpusets)

```
* For details of cpumask_scnprintf() and cpumask_parse(),
* see bitmap_scnprintf() and bitmap_parse() in lib/bitmap.c.
+ * For details of cpulist_scnprintf() and cpulist_parse(), see
+ * bitmap_scnlistprintf() and bitmap_parselist(), also in bitmap.c.
*
* The available cpumask operations are:
*
@@ -46,6 +48,8 @@
*
* int cpumask_scnprintf(buf, len, mask) Format cpumask for printing
* int cpumask_parse(ubuf, ulen, mask) Parse ascii string as cpumask
+ * int cpulist_scnprintf(buf, len, mask) Format cpumask as list for printing
+ * int cpulist_parse(buf, map) Parse ascii string as cpulist
*
* for_each_cpu_mask(cpu, mask) for-loop cpu over mask
*
@@ -268,14 +272,28 @@ static inline int __cpumask_scnprintf(ch
    return bitmap_scnprintf(buf, len, srcp->bits, nbits);
}

-#define cpumask_parse(ubuf, ulen, src) \
- __cpumask_parse((ubuf), (ulen), &(src), NR_CPUS)
+#define cpumask_parse(ubuf, ulen, dst) \
+ __cpumask_parse((ubuf), (ulen), &(dst), NR_CPUS)
static inline int __cpumask_parse(const char __user *buf, int len,
                                cpumask_t *dstp, int nbits)
{
    return bitmap_parse(buf, len, dstp->bits, nbits);
}

+#define cpulist_scnprintf(buf, len, src) \
+ __cpulist_scnprintf((buf), (len), &(src), NR_CPUS)
+static inline int __cpulist_scnprintf(char *buf, int len,
+ const cpumask_t *srcp, int nbits)
+{
+ return bitmap_scnlistprintf(buf, len, srcp->bits, nbits);
+}
+
+#define cpulist_parse(buf, dst) __cpulist_parse((buf), &(dst), NR_CPUS)
+static inline int __cpulist_parse(const char *buf, cpumask_t *dstp, int nbits)
+{
+ return bitmap_parselist(buf, dstp->bits, nbits);
+}
+
+#if NR_CPUS > 1
#define for_each_cpu_mask(cpu, mask) \
    for ((cpu) = first_cpu(mask); \
Index: 2.6.8-rc2-mm2/include/linux/nodemask.h
=====
--- 2.6.8-rc2-mm2.orig/include/linux/nodemask.h 2004-08-08 23:17:35.000000000 -0700
+++ 2.6.8-rc2-mm2/include/linux/nodemask.h 2004-08-08 23:24:57.000000000 -0700
```

Linux-Kernel: Re: [Lse-tech] [PATCH] new bitmap list format (for cpusets)

@@ -10,6 +10,8 @@

```
*
* For details of nodemask_scnprintf() and nodemask_parse(),
* see bitmap_scnprintf() and bitmap_parse() in lib/bitmap.c.
+ * For details of nodelist_scnprintf() and nodelist_parse(), see
+ * bitmap_scnlistprintf() and bitmap_parselist(), also in bitmap.c.
```

```
*
* The available nodemask operations are:
```

@@ -46,6 +48,8 @@

```
*
* int nodemask_scnprintf(buf, len, mask) Format nodemask for printing
* int nodemask_parse(ubuf, ulen, mask) Parse ascii string as nodemask
+ * int nodelist_scnprintf(buf, len, mask) Format nodemask as list for printing
+ * int nodelist_parse(buf, map) Parse ascii string as nodelist
```

```
*
* for_each_node_mask(node, mask) for-loop node over mask
```

```
@@ -271,14 +275,28 @@ static inline int __nodemask_scnprintf(c
    return bitmap_scnprintf(buf, len, srcp->bits, nbits);
}
```

```

-#define nodemask_parse(ubuf, ulen, src) \
- __nodemask_parse((ubuf), (ulen), &(src), MAX_NUMNODES)
+#define nodemask_parse(ubuf, ulen, dst) \
+ __nodemask_parse((ubuf), (ulen), &(dst), MAX_NUMNODES)
static inline int __nodemask_parse(const char __user *buf, int len,
    nodemask_t *dstp, int nbits)
{
    return bitmap_parse(buf, len, dstp->bits, nbits);
}
```

```

+#define nodelist_scnprintf(buf, len, src) \
+ __nodelist_scnprintf((buf), (len), &(src), MAX_NUMNODES)
+static inline int __nodelist_scnprintf(char *buf, int len,
+ const nodemask_t *srcp, int nbits)
+{
+ return bitmap_scnlistprintf(buf, len, srcp->bits, nbits);
+}
+
+#define nodelist_parse(buf, dst) __nodelist_parse((buf), &(dst), MAX_NUMNODES)
+static inline int __nodelist_parse(const char *buf, nodemask_t *dstp, int nbits)
+{
+ return bitmap_parselist(buf, dstp->bits, nbits);
+}
```

```

+
+#if MAX_NUMNODES > 1
+#define for_each_node_mask(node, mask) \
    for ((node) = first_node(mask); \
```

Index: 2.6.8-rc2-mm2/lib/bitmap.c

Linux-Kernel: Re: [Lse-tech] [PATCH] new bitmap list format (for cpusets)

```
--- 2.6.8-rc2-mm2.orig/lib/bitmap.c 2004-08-08 23:17:35.000000000 -0700
+++ 2.6.8-rc2-mm2/lib/bitmap.c 2004-08-09 00:11:57.000000000 -0700
@@ -291,6 +291,7 @@ EXPORT_SYMBOL(__bitmap_weight);
#define nbits_to_hold_value(val) fls(val)
#define roundup_power2(val,modulus) (((val) + (modulus) - 1) & ~((modulus) - 1))
#define unhex(c) (isdigit(c) ? (c - '0') : (toupper(c) - 'A' + 10))
+#define BASEDEC 10 /* fancier cpuset lists input in decimal */

/**
 * bitmap_scnprintf - convert bitmap to an ASCII hex string.
@@ -409,6 +410,108 @@ int bitmap_parse(const char __user *ubuf
}
EXPORT_SYMBOL(bitmap_parse);

+/*
+ * bscnl_emit(buf, buflen, rbot, rtop, bp)
+ *
+ * Helper routine for bitmap_scnlistprintf(). Write decimal number
+ * or range to buf, suppressing output past buf+buflen, with optional
+ * comma-prefix. Return len of what would be written to buf, if it
+ * all fit.
+ */
+static inline int bscnl_emit(char *buf, int buflen, int rbot, int rtop, int len)
+{
+ if (len > 0)
+ len += scnprintf(buf + len, buflen - len, ",");
+ if (rbot == rtop)
+ len += scnprintf(buf + len, buflen - len, "%d", rbot);
+ else
+ len += scnprintf(buf + len, buflen - len, "%d-%d", rbot, rtop);
+ return len;
+}
+
+/**
+ * bitmap_scnlistprintf - convert bitmap to list format ASCII string
+ * @buf: byte buffer into which string is placed
+ * @buflen: reserved size of @buf, in bytes
+ * @maskp: pointer to bitmap to convert
+ * @nmaskbits: size of bitmap, in bits
+ *
+ * Output format is a comma-separated list of decimal numbers and
+ * ranges. Consecutively set bits are shown as two hyphen-separated
+ * decimal numbers, the smallest and largest bit numbers set in
+ * the range. Output format is compatible with the format
+ * accepted as input by bitmap_parselist().
+ *
+ * The return value is the number of characters which would be
+ * generated for the given input, excluding the trailing '\0', as
+ * per ISO C99.
+ */
+int bitmap_scnlistprintf(char *buf, unsigned int buflen,
```

Linux-Kernel: Re: [Lse-tech] [PATCH] new bitmap list format (for cpusets)

```
+ const unsigned long *maskp, int nmaskbits)
+{
+ int len = 0;
+ /* current bit is 'cur', most recently seen range is [rbot, rtop] */
+ int cur, rbot, rtop;
+
+ rbot = cur = find_first_bit(maskp, nmaskbits);
+ while (cur < nmaskbits) {
+ rtop = cur;
+ cur = find_next_bit(maskp, nmaskbits, cur+1);
+ if (cur >= nmaskbits || cur > rtop + 1) {
+ len = bscnl_emit(buf, buflen, rbot, rtop, len);
+ rbot = cur;
+ }
+ }
+ return len;
+}
+EXPORT_SYMBOL(bitmap_scnlistprintf);
+
+/**
+ * bitmap_parselist – convert list format ASCII string to bitmap
+ * @buf: read nul-terminated user string from this buffer
+ * @mask: write resulting mask here
+ * @nmaskbits: number of bits in mask to be written
+ *
+ * Input format is a comma-separated list of decimal numbers and
+ * ranges. Consecutively set bits are shown as two hyphen-separated
+ * decimal numbers, the smallest and largest bit numbers set in
+ * the range.
+ *
+ * Returns 0 on success, -errno on invalid input strings:
+ * -EINVAL: second number in range smaller than first
+ * -EINVAL: invalid character in string
+ * -ERANGE: bit number specified too large for mask
+ */
+int bitmap_parselist(const char *bp, unsigned long *maskp, int nmaskbits)
+{
+ unsigned a, b;
+
+ bitmap_zero(maskp, nmaskbits);
+ do {
+ if (!isdigit(*bp))
+ return -EINVAL;
+ b = a = simple_strtoul(bp, (char **)&bp, BASEDEC);
+ if (*bp == '-') {
+ bp++;
+ if (!isdigit(*bp))
+ return -EINVAL;
+ b = simple_strtoul(bp, (char **)&bp, BASEDEC);
+ }
+ if (!(a <= b))
```

Linux-Kernel: Re: [Lse-tech] [PATCH] new bitmap list format (for cpusets)

```
+ return -EINVAL;
+ if (b >= nmaskbits)
+ return -ERANGE;
+ while (a <= b) {
+ set_bit(a, maskp);
+ a++;
+ }
+ if (*bp == ',')
+ bp++;
+ } while (*bp != '\0' && *bp != '\n');
+ return 0;
+}
+EXPORT_SYMBOL(bitmap_parselist);
+
/**
 * bitmap_find_free_region - find a contiguous aligned mem region
 * @bitmap: an array of unsigned longs corresponding to the bitmap
```

--

```
      I won't rest till it's the best ...
      Programmer, Linux Scalability
      Paul Jackson <pj@sgi.com> 1.650.933.1373
```

-

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@vger.kernel.org
More majordomo info at <http://vger.kernel.org/majordomo-info.html>
Please read the FAQ at <http://www.tux.org/lkml/>