Subject: Re: [PATCH 16/28] [FLAT 1/6] Changes in data structures for flat model
Posted by Sukadev Bhattiprolu on Tue, 19 Jun 2007 06:52:35 GMT
View Forum Message <> Reply to Message

Pavel Emelianov [xemul@openvz.org] wrote:
| sukadev@us.ibm.com wrote:
| > Pavel Emelianov [xemul@openvz.org] wrote:
| > | This patch opens the flat model patches.
| > |
| > | The flat model idea is that struct pid has two numbers. The first one
| > | (pid->nr) is a global one and is unique in the system. The second one
| > | (pid->vnr) is a virtual pid. It is used on the kernel user boundary only.
| >
| > This approach duplicates 5 integers and 2 pointers per process for every
| > process in the system. While this may not be expensive for processes that
| > actually use multiple namespaces, doesn't it waste memory if majority of
| > processes exist only in one namespace ?
|
| task_struct alignment allows for it. so does the alignment of signal structure.
| and please note that this comes with appropriate ifdefs around. the only problem
| is with struct pid, but we're virtualizing it after all!

Hmm. I don't understand the last part "we are virtualizing 'struct pid'".
Even so, with the FLAT model, every process will still have two
pid_t values, two hash-chain links etc - no ?

|
| moreover - two integers and a pointer to the namespace is the minimal set of
| fields for pid that is visible from two namespaces...

I ignored the pid_namespace pointer. But even a process that exists only
in init_pid_ns would have the extra fields right ?