

---

Subject: Re: [PATCH v2 11/11] protect architectures where THREAD\_SIZE >= PAGE\_SIZE against fork bombs

Posted by [KAMEZAWA Hiroyuki](#) on Fri, 10 Aug 2012 17:54:48 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

(2012/08/09 22:01), Glauber Costa wrote:

> Because those architectures will draw their stacks directly from the  
> page allocator, rather than the slab cache, we can directly pass  
> \_\_GFP\_KMEMCG flag, and issue the corresponding free\_pages.  
>  
> This code path is taken when the architecture doesn't define  
> CONFIG\_ARCH\_THREAD\_INFO\_ALLOCATOR (only ia64 seems to), and has  
> THREAD\_SIZE >= PAGE\_SIZE. Luckily, most - if not all - of the remaining  
> architectures fall in this category.  
>  
> This will guarantee that every stack page is accounted to the memcg the  
> process currently lives on, and will have the allocations to fail if  
> they go over limit.  
>  
> For the time being, I am defining a new variant of THREADINFO\_GFP, not  
> to mess with the other path. Once the slab is also tracked by memcg, we  
> can get rid of that flag.  
>  
> Tested to successfully protect against :(){ :|:& };:  
>  
> Signed-off-by: Glauber Costa <glommer@parallels.com>  
> Acked-by: Frederic Weisbecker <fweisbec@redhat.com>  
> CC: Christoph Lameter <cl@linux.com>  
> CC: Pekka Enberg <penberg@cs.helsinki.fi>  
> CC: Michal Hocko <mhocko@suse.cz>  
> CC: Kamezawa Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com>  
> CC: Johannes Weiner <hannes@cmpxchg.org>  
> CC: Suleiman Souhlal <suleiman@google.com>

Acked-by: KAMEZAWA Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com>

> ---  
> include/linux/thread\_info.h | 2 ++  
> kernel/fork.c | 4 +++-  
> 2 files changed, 4 insertions(+), 2 deletions(-)  
>  
> diff --git a/include/linux/thread\_info.h b/include/linux/thread\_info.h  
> index ccc1899..e7e0473 100644  
> --- a/include/linux/thread\_info.h  
> +++ b/include/linux/thread\_info.h  
> @@ -61,6 +61,8 @@ extern long do\_no\_restart\_syscall(struct restart\_block \*parm);  
> # define THREADINFO\_GFP (GFP\_KERNEL | \_\_GFP\_NOTRACK)

```
> #endif
>
> +#define THREADINFO_GFP_ACCOUNTED (THREADINFO_GFP | __GFP_KMEMCG)
> +
> /*
>   * flag set/clear/test wrappers
>   * - pass TIF_xxxx constants to these functions
> diff --git a/kernel/fork.c b/kernel/fork.c
> index dc3ff16..b0b90c3 100644
> --- a/kernel/fork.c
> +++ b/kernel/fork.c
> @@ -142,7 +142,7 @@ void __weak arch_release_thread_info(struct thread_info *ti) { }
> static struct thread_info *alloc_thread_info_node(struct task_struct *tsk,
>         int node)
> {
> - struct page *page = alloc_pages_node(node, THREADINFO_GFP,
> + struct page *page = alloc_pages_node(node, THREADINFO_GFP_ACCOUNTED,
>         THREAD_SIZE_ORDER);
>
>     return page ? page_address(page) : NULL;
> @@ -151,7 +151,7 @@ static struct thread_info *alloc_thread_info_node(struct task_struct
*tsk,
>     static inline void free_thread_info(struct thread_info *ti)
> {
>     arch_release_thread_info(ti);
> - free_pages((unsigned long)ti, THREAD_SIZE_ORDER);
> + free_accounted_pages((unsigned long)ti, THREAD_SIZE_ORDER);
> }
> # else
> static struct kmem_cache *thread_info_cache;
>
```

---