

---

Subject: Re: [PATCH 3/3] add the clone64() and unshare64() syscalls

Posted by [hpa](#) on Thu, 10 Apr 2008 03:40:07 GMT

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

---

sukadev@us.ibm.com wrote:

> Jakub Jelinek [jakub@redhat.com] wrote:

> | On Wed, Apr 09, 2008 at 03:34:59PM -0700, sukadev@us.ibm.com wrote:

> | > From: Cedric Le Goater <clg@fr.ibm.com>

> | > Subject: [PATCH 3/3] add the clone64() and unshare64() syscalls

> | >

> | > This patch adds 2 new syscalls :

> | >

> | > long sys\_clone64(unsigned long flags\_high, unsigned long flags\_low,

> | > unsigned long newsp);

> | >

> | > long sys\_unshare64(unsigned long flags\_high, unsigned long flags\_low);

> |

> | Can you explain why are you adding it for 64-bit arches too? unsigned long

> | is there already 64-bit, and both sys\_clone and sys\_unshare have unsigned

> | long flags, rather than unsigned int.

>

> Hmm,

>

> By simply resuing clone() on 64 bit and adding a new call for 32-bit won't

> the semantics of clone() differ between the two ?

>

> i.e clone() on 64 bit supports say CLONE\_NEWPTS clone() on 32bit does not ?

>

> Wouldn't it be simpler/cleaner if clone() and clone64() behaved the same

> on both 32 and 64 bit systems ?

>

No, not really. The way this work on the libc side is pretty much "use clone64 if it exists, otherwise use clone".

-hpa

---

Containers mailing list

[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)

<https://lists.linux-foundation.org/mailman/listinfo/containers>

---