
Subject: [RFC PATCH] posix timers: allocate timer id per task
Posted by [Stanislav Kinsbursky](#) on Mon, 15 Oct 2012 16:17:29 GMT
[View Forum Message](#) <> [Reply to Message](#)

This patch is required CRIU project (www.criu.org).
To migrate processes with posix timers we have to make sure, that we can restore posix timer with proper id.
Currently, this is not true, because timer ids are allocated globally.
So, this is precursor patch and it's purpose is make posix timer id to be allocated per task.

Patch replaces global idr with global hash table for posix timers and makes timer ids unique not globally, but per task. Next free timer id is type of integer and stored on signal struct (posix_timer_id). If free timer id reaches negative value on timer creation, it will be dropped to zero and -EAGAIN will be returned to user.

Hash table is size of page (4KB).

Key is constructed as follows:

```
key = hash_ptr(current->signal) ^ hash_32(posix_timer_id);
```

Signed-off-by: Stanislav Kinsbursky <skinsbursky@parallels.com>

```
---
include/linux/posix-timers.h | 1
include/linux/sched.h        | 4 +-
kernel/posix-timers.c        | 111 ++++++-----
3 files changed, 77 insertions(+), 39 deletions(-)
```

```
diff --git a/include/linux/posix-timers.h b/include/linux/posix-timers.h
```

```
index 042058f..60bac69 100644
```

```
--- a/include/linux/posix-timers.h
```

```
+++ b/include/linux/posix-timers.h
```

```
@ @ -55,6 +55,7 @ @ struct cpu_timer_list {
```

```
/* POSIX.1b interval timer structure. */
```

```
struct k_itimer {
```

```
    struct list_head list; /* free/ allocate list */
```

```
+ struct hlist_node t_hash;
```

```
    spinlock_t it_lock;
```

```
    clockid_t it_clock; /* which timer type */
```

```
    timer_t it_id; /* timer id */
```

```
diff --git a/include/linux/sched.h b/include/linux/sched.h
```

```
index 0dd42a0..dce1651 100644
```

```
--- a/include/linux/sched.h
```

```
+++ b/include/linux/sched.h
```

```
@ @ -51,6 +51,7 @ @ struct sched_param {
```

```
#include <linux/cred.h>
```

```
#include <linux/llist.h>
```

```
#include <linux/uidgid.h>
```

```

#include <linux/idr.h>

#include <asm/processor.h>

@@ -536,7 +537,8 @@ struct signal_struct {
    unsigned int  has_child_subreaper:1;

    /* POSIX.1b Interval Timers */
- struct list_head posix_timers;
+ int  posix_timer_id;
+ struct list_head posix_timers;

    /* ITIMER_REAL timer for the process */
    struct hrtimer real_timer;
diff --git a/kernel/posix-timers.c b/kernel/posix-timers.c
index 69185ae..4355a9d 100644
--- a/kernel/posix-timers.c
+++ b/kernel/posix-timers.c
@@ -47,31 +47,27 @@
#include <linux/wait.h>
#include <linux/workqueue.h>
#include <linux/export.h>
#include <linux/hash.h>

/*
- * Management arrays for POSIX timers. Timers are kept in slab memory
- * Timer ids are allocated by an external routine that keeps track of the
- * id and the timer. The external interface is:
- *
- * void *idr_find(struct idr *idp, int id);      to find timer_id <id>
- * int idr_get_new(struct idr *idp, void *ptr);  to get a new id and
- *                                             related it to <ptr>
- * void idr_remove(struct idr *idp, int id);     to release <id>
- * void idr_init(struct idr *idp);              to initialize <idp>
- *                                             which we supply.
- * The idr_get_new *may* call slab for more memory so it must not be
- * called under a spin lock. Likewise idr_remove may release memory
- * (but it may be ok to do this under a lock...).
- * idr_find is just a memory look up and is quite fast. A -1 return
- * indicates that the requested id does not exist.
+ * Management arrays for POSIX timers. Timers are now kept in static PAGE-size
+ * hash table.
+ * Timer ids are allocated by local routine, which selects proper hash head by
+ * key, constructed from current->signal address and per signal struct counter.
+ * This keeps timer ids unique per task, but now they can intersect between
+ * tasks.
 */

```

```

/*
 * Lets keep our timers in a slab cache :-)
 */
static struct kmem_cache *posix_timers_cache;
-static struct idr posix_timers_idr;
-static DEFINE_SPINLOCK(idr_lock);
+
+#define POSIX_TIMERS_HASH_SIZE (PAGE_SIZE / sizeof(struct hlist_head))
+#define POSIX_TIMERS_HASH_BITS 9
+
+static struct hlist_head posix_timers_hashtable[POSIX_TIMERS_HASH_SIZE];
+static DEFINE_SPINLOCK(hash_lock);

/*
 * we assume that the new SIGEV_THREAD_ID shares no bits with the other
 @@ -152,6 +148,56 @@ static struct k_itimer *__lock_timer(timer_t timer_id, unsigned long
 *flags);
 __timr;      \
})

+static int hash(struct signal_struct *sig, unsigned int nr)
+{
+ int hash = hash_ptr(sig, POSIX_TIMERS_HASH_BITS);
+ return hash ^ hash_32(nr, POSIX_TIMERS_HASH_BITS);
+}
+
+static struct k_itimer *__posix_timers_find(struct hlist_head *head, struct signal_struct *sig,
timer_t id)
+{
+ struct hlist_node *node;
+ struct k_itimer *timer;
+
+ hlist_for_each_entry(timer, node, head, t_hash) {
+ if ((timer->it_signal == sig) && (timer->it_id == id))
+ return timer;
+ }
+ return NULL;
+}
+
+static struct k_itimer *posix_timer_find(timer_t id, unsigned long *flags)
+{
+ struct k_itimer *timer;
+ struct signal_struct *sig = current->signal;
+ struct hlist_head *head = &posix_timers_hashtable[hash(sig, id)];
+
+ spin_lock_irqsave(&hash_lock, *flags);
+ timer = __posix_timers_find(head, sig, id);
+ spin_unlock_irqrestore(&hash_lock, *flags);

```

```

+ return timer;
+}
+
+static int posix_timer_add(struct k_itimer *timer)
+{
+ struct signal_struct *sig = current->signal;
+ int next_free_id = sig->posix_timer_id;
+ struct hlist_head *head;
+ int ret = -ENOENT;
+
+ do {
+ spin_lock(&hash_lock);
+ head = &posix_timers_hashtable[hash(sig, sig->posix_timer_id)];
+ if (__posix_timers_find(head, sig, sig->posix_timer_id) == NULL) {
+ hlist_add_head(&timer->t_hash, head);
+ ret = sig->posix_timer_id++;
+ } else if (++sig->posix_timer_id == next_free_id)
+ ret = -EAGAIN;
+ spin_unlock(&hash_lock);
+ } while (ret == -ENOENT);
+ return ret;
+}
+
+static inline void unlock_timer(struct k_itimer *timr, unsigned long flags)
+{
+ spin_unlock_irqrestore(&timr->it_lock, flags);
@@ -271,6 +317,7 @@ static __init int init_posix_timers(void)
+ .timer_get = common_timer_get,
+ .timer_del = common_timer_del,
+ };
+ int i;

+ posix_timers_register_clock(CLOCK_REALTIME, &clock_realtime);
+ posix_timers_register_clock(CLOCK_MONOTONIC, &clock_monotonic);
@@ -282,7 +329,8 @@ static __init int init_posix_timers(void)
+ posix_timers_cache = kmem_cache_create("posix_timers_cache",
+ sizeof (struct k_itimer), 0, SLAB_PANIC,
+ NULL);
+ idr_init(&posix_timers_id);
+ for (i = 0; i < POSIX_TIMERS_HASH_SIZE; i++)
+ INIT_HLIST_HEAD(&posix_timers_hashtable[i]);
+ return 0;
+ }

@@ -504,9 +552,9 @@ static void release_posix_timer(struct k_itimer *tmr, int it_id_set)
+ {
+ if (it_id_set) {
+ unsigned long flags;

```

```

- spin_lock_irqsave(&idr_lock, flags);
- idr_remove(&posix_timers_id, tmr->it_id);
- spin_unlock_irqrestore(&idr_lock, flags);
+ spin_lock_irqsave(&hash_lock, flags);
+ hlist_del(&tmr->t_hash);
+ spin_unlock_irqrestore(&hash_lock, flags);
}
put_pid(tmr->it_pid);
sigqueue_free(tmr->sigq);
@@ -552,22 +600,9 @@ SYSCALL_DEFINE3(timer_create, const clockid_t, which_clock,
    return -EAGAIN;

    spin_lock_init(&new_timer->it_lock);
- retry:
- if (unlikely(!idr_pre_get(&posix_timers_id, GFP_KERNEL))) {
-     error = -EAGAIN;
-     goto out;
- }
- spin_lock_irq(&idr_lock);
- error = idr_get_new(&posix_timers_id, new_timer, &new_timer_id);
- spin_unlock_irq(&idr_lock);
- if (error) {
-     if (error == -EAGAIN)
-         goto retry;
-     /*
-      * Weird looking, but we return EAGAIN if the IDR is
-      * full (proper POSIX return value for this)
-      */
-     error = -EAGAIN;
+ new_timer_id = posix_timer_add(new_timer);
+ if (new_timer_id < 0) {
+     error = new_timer_id;
+     goto out;
+ }

@@ -640,7 +675,7 @@ static struct k_itimer *__lock_timer(timer_t timer_id, unsigned long *flags)
    struct k_itimer *timr;

    rcu_read_lock();
- timr = idr_find(&posix_timers_id, (int)timer_id);
+ timr = posix_timer_find(timer_id, flags);
    if (timr) {
        spin_lock_irqsave(&timr->it_lock, *flags);
        if (timr->it_signal == current->signal) {

```

Subject: Re: [RFC PATCH] posix timers: allocate timer id per task

Posted by [Eric Dumazet](#) on Mon, 15 Oct 2012 16:34:43 GMT

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

On Mon, 2012-10-15 at 20:17 +0400, Stanislav Kinsbursky wrote:

- > This patch is required CRIU project (www.criu.org).
- > To migrate processes with posix timers we have to make sure, that we can
- > restore posix timer with proper id.
- > Currently, this is not true, because timer ids are allocated globally.
- > So, this is precursor patch and it's purpose is make posix timer id to be
- > allocated per task.
- >
- > Patch replaces global idr with global hash table for posix timers and
- > makes timer ids unique not globally, but per task. Next free timer id is type
- > of integer and stored on signal struct (posix_timer_id). If free timer id
- > reaches negative value on timer creation, it will be dropped to zero and
- > -EAGAIN will be returned to user.
- > Hash table is size of page (4KB).
- > Key is constructed as follows:
- > key = hash_ptr(current->signal) ^ hash_32(posix_timer_id);
- >
- > Signed-off-by: Stanislav Kinsbursky <skinsbursky@parallels.com>

Hmm, it seems you removed idr, rcu friendly, and reinstated a fixed size hash table, protected by a `_single_` spinlock ? Oh well.

Please take a look at commit `8af088710d1e`, and make sure you fix your problem and keep performance as good as before.

Subject: Re: [RFC PATCH] posix timers: allocate timer id per task

Posted by [Peter Zijlstra](#) on Mon, 15 Oct 2012 17:04:05 GMT

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

On Mon, 2012-10-15 at 20:17 +0400, Stanislav Kinsbursky wrote:

- > This patch is required CRIU project (www.criu.org).
- > To migrate processes with posix timers we have to make sure, that we can
- > restore posix timer with proper id.
- > Currently, this is not true, because timer ids are allocated globally.
- > So, this is precursor patch and it's purpose is make posix timer id to be
- > allocated per task.
- >
- > Patch replaces global idr with global hash table for posix timers and
- > makes timer ids unique not globally, but per task. Next free timer id is type
- > of integer and stored on signal struct (posix_timer_id). If free timer id
- > reaches negative value on timer creation, it will be dropped to zero and
- > -EAGAIN will be returned to user.
- > Hash table is size of page (4KB).

> Key is constructed as follows:

> key = hash_ptr(current->signal) ^ hash_32(posix_timer_id);

but but but.. isn't this what namespaces were invented for to solve? Why not use the regular namespace infrastructure?

Subject: Re: [RFC PATCH] posix timers: allocate timer id per task

Posted by [Thomas Gleixner](#) on Mon, 15 Oct 2012 19:08:18 GMT

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

On Mon, 15 Oct 2012, Stanislav Kinsbursky wrote:

> This patch is required CRIU project (www.criu.org).

> To migrate processes with posix timers we have to make sure, that we can
> restore posix timer with proper id.

> Currently, this is not true, because timer ids are allocated globally.

> So, this is precursor patch and it's purpose is make posix timer id to be
> allocated per task.

You can't allocate them per task. posix timers are process wide.

What's the reason why you did not make the posix timer ids per name space instead of going down to the per process level ?

> Patch replaces global idr with global hash table for posix timers and

> makes timer ids unique not globally, but per task. Next free timer id is type

> of integer and stored on signal struct (posix_timer_id). If free timer id

> reaches negative value on timer creation, it will be dropped to zero and

> -EAGAIN will be returned to user.

So you want to allow 2^{31} posix timers created for a single process?

> +static struct k_itimer *__posix_timers_find(struct hlist_head *head, struct signal_struct *sig,
timer_t id)

> +{

> + struct hlist_node *node;

> + struct k_itimer *timer;

> +

> + hlist_for_each_entry(timer, node, head, t_hash) {

> + if ((timer->it_signal == sig) && (timer->it_id == id))

> + return timer;

> + }

> + return NULL;

> + }

> +

> +static struct k_itimer *posix_timer_find(timer_t id, unsigned long *flags)

> +{

```
> + struct k_itimer *timer;
> + struct signal_struct *sig = current->signal;
> + struct hlist_head *head = &posix_timers_hashtable[hash(sig, id)];
> +
> + spin_lock_irqsave(&hash_lock, *flags);
```

This is not going to fly. You just reintroduced a massive scalability problem. See commit 8af08871

Thanks,

tglx

Subject: Re: [RFC PATCH] posix timers: allocate timer id per task
Posted by [Stanislav Kinsbursky](#) on Tue, 16 Oct 2012 07:57:58 GMT
[View Forum Message](#) <> [Reply to Message](#)

```
> On Mon, 2012-10-15 at 20:17 +0400, Stanislav Kinsbursky wrote:
>> This patch is required CRIU project (www.criu.org).
>> To migrate processes with posix timers we have to make sure, that we can
>> restore posix timer with proper id.
>> Currently, this is not true, because timer ids are allocated globally.
>> So, this is precursor patch and it's purpose is make posix timer id to be
>> allocated per task.
>>
>> Patch replaces global idr with global hash table for posix timers and
>> makes timer ids unique not globally, but per task. Next free timer id is type
>> of integer and stored on signal struct (posix_timer_id). If free timer id
>> reaches negative value on timer creation, it will be dropped to zero and
>> -EAGAIN will be returned to user.
>> Hash table is size of page (4KB).
>> Key is constructed as follows:
>> key = hash_ptr(current->signal) ^ hash_32(posix_timer_id);
>>
>> Signed-off-by: Stanislav Kinsbursky <skinsbursky@parallels.com>
>
>
> Hmm, it seems you removed idr, rcu friendly, and reinstated a fixed size
> hash table, protected by a _single_ spinlock ? Oh well.
>
> Please take a look at commit 8af088710d1e, and make sure you fix your
> problem and keep performance as good as before.
>
```

Thanks, Eric.
I'll update.

>
>

--

Best regards,
Stanislav Kinsbursky

Subject: Re: [RFC PATCH] posix timers: allocate timer id per task
Posted by [Stanislav Kinsbursky](#) on Tue, 16 Oct 2012 08:00:06 GMT
[View Forum Message](#) <> [Reply to Message](#)

> On Mon, 2012-10-15 at 20:17 +0400, Stanislav Kinsbursky wrote:
>> This patch is required CRIU project (www.criu.org).
>> To migrate processes with posix timers we have to make sure, that we can
>> restore posix timer with proper id.
>> Currently, this is not true, because timer ids are allocated globally.
>> So, this is precursor patch and it's purpose is make posix timer id to be
>> allocated per task.
>>
>> Patch replaces global idr with global hash table for posix timers and
>> makes timer ids unique not globally, but per task. Next free timer id is type
>> of integer and stored on signal struct (posix_timer_id). If free timer id
>> reaches negative value on timer creation, it will be dropped to zero and
>> -EAGAIN will be returned to user.
>> Hash table is size of page (4KB).
>> Key is constructed as follows:
>> key = hash_ptr(current->signal) ^ hash_32(posix_timer_id);
>
> but but but.. isn't this what namespaces were invented for to solve? Why
> not use the regular namespace infrastructure?
>

The reason is that CRIU have to support single processes within existent namespaces.

--

Best regards,
Stanislav Kinsbursky

Subject: Re: [RFC PATCH] posix timers: allocate timer id per task
Posted by [Stanislav Kinsbursky](#) on Tue, 16 Oct 2012 08:08:14 GMT
[View Forum Message](#) <> [Reply to Message](#)

> On Mon, 15 Oct 2012, Stanislav Kinsbursky wrote:
>
>> This patch is required CRIU project (www.criu.org).
>> To migrate processes with posix timers we have to make sure, that we can
>> restore posix timer with proper id.
>> Currently, this is not true, because timer ids are allocated globally.
>> So, this is precursor patch and it's purpose is make posix timer id to be
>> allocated per task.
>
> You can't allocate them per task. posix timers are process wide.
>

This is probably a misunderstanding.
I meant process process.

> What's the reason why you did not make the posix timer ids per name
> space instead of going down to the per process level ?
>

The reason is that CRIU have to support single processes regardless to namespaces.

>> Patch replaces global idr with global hash table for posix timers and
>> makes timer ids unique not globally, but per task. Next free timer id is type
>> of integer and stored on signal struct (posix_timer_id). If free timer id
>> reaches negative value on timer creation, it will be dropped to zero and
>> -EAGAIN will be returned to user.
>
> So you want to allow 2^{31} posix timers created for a single process?
>

I don't really want it.
I just tried to preserve existent logic. I.e. currently one process can create
 2^{31} posix timers (if no other processes tried to do the same).
But the amount can be decreased, of course.

```
>> +static struct k_itimer *__posix_timers_find(struct hlist_head *head, struct signal_struct *sig,
timer_t id)
>> +{
>> + struct hlist_node *node;
>> + struct k_itimer *timer;
>> +
>> + hlist_for_each_entry(timer, node, head, t_hash) {
>> + if ((timer->it_signal == sig) && (timer->it_id == id))
>> + return timer;
>> + }
>> + return NULL;
>> +}
```

```
>> +
>> +static struct k_itimer *posix_timer_find(timer_t id, unsigned long *flags)
>> +{
>> + struct k_itimer *timer;
>> + struct signal_struct *sig = current->signal;
>> + struct hlist_head *head = &posix_timers_hashtable[hash(sig, id)];
>> +
>> + spin_lock_irqsave(&hash_lock, *flags);
>
> This is not going to fly. You just reintroduced a massive scalability
> problem. See commit 8af08871
>
```

Yep, Eric already pointed to it.
I'll try to fix this problem, if the idea with hash table suits in general.
Thanks.

> Thanks,
>
> tglx
>

--

Best regards,
Stanislav Kinsbursky
