
Subject: [RFC PATCH 1/1] SYSVIPC/semaphores: document the new per <pid> semundo proc file

Posted by [Nadia Derby](#) on Thu, 26 Jun 2008 11:34:58 GMT

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

PATCH [01/01]

This patch completes the patch series I've sent yesterday to introduce the new per process semundo proc file.

(see <https://lists.linux-foundation.org/pipermail/containers/2008-June/011502.html>)

It updates Documentation/filesystems/proc.txt with that file's description.

Please, tell me if I need to resend the complete series.

Regards,
Nadia

Signed-off-by: Nadia Derby <Nadia.Derbey@bull.net>

Documentation/filesystems/proc.txt | 56 ++++++-----
1 file changed, 54 insertions(+), 2 deletions(-)

Index: linux-2.6.26-rc5-mm3.4.semundo/Documentation/filesystems/proc.txt

=====

--- linux-2.6.26-rc5-mm3.4.semundo.orig/Documentation/filesystems/proc.txt 2008-06-24 09:05:00.000000000 +0200

+++ linux-2.6.26-rc5-mm3.4.semundo/Documentation/filesystems/proc.txt 2008-06-26 14:35:27.000000000 +0200

@ @ -44,7 +44,8 @ @ Table of Contents

- 2.14 /proc/<pid>/io - Display the IO accounting fields
- 2.15 /proc/<pid>/coredump_filter - Core dump filtering settings
- 2.16 /proc/<pid>/mountinfo - Information about mounts
- 2.17 /proc/sys/fs/types - File system type specific parameters
- + 2.17 /proc/<pid>/semundo - Information sysv semaphores undo list
- + 2.18 /proc/sys/fs/types - File system type specific parameters

Preface

@ @ -2453,7 +2454,58 @ @ For more information on mount propagatio

Documentation/filesystems/sharedsubtree.txt

-2.17 /proc/sys/fs/types/ - File system type specific parameters

+2.17 /proc/<pid>/semundo - Information about sysv semaphores undo list

+-----

+

+This file contains lines of the form:

```

+
+294914    1   -1    1    1    1    1    1    1    1
+ (1)                (i)
+
+(1) semaphore id: semaphore set identifier as returned by a semget()
+(i) adj value: undo count for the i-th semaphore in the semaphore set.
+      This is the value that will be added to the i-th semaphore when
+      the process <pid> exits.
+
+Ex: when the process owning the semundo procfs file given above exits, the
+semaphore set whose id is 294914 will have its semaphores values updated as
+follows:
+ . semaphore # 0 will have its value incremented by 1
+ . semaphore # 1 will have its value decremented by 1
+ . semaphore # 2 will have its value incremented by 1, etc.
+
+This file is particularly useful if one wants to take a snapshot of a process'
+state in terms of sysv ipc's, during a checkpoint phase, for example.
+
+This file can be open for reading by any process.
+It can also be open for writing, but the write operation is only allowed into
+/proc/self/semundo.
+
+The write syntax is the following:
+<semaphore set id> <VAL0> ... <VALi> ... <VALn>
+
+with <VALi> = new undo value for semaphore # <i> in the semaphore set.
+
+Note that the number of values given after the semaphore id should be
+identical to the number of semaphores in the set (i.e. to the nsems parameter
+given in the semget() syscall).
+
+Writing to the semundo file has the following effect:
+ . if the owning process already had an undo list, each semaphore in the set
+   will have its undo count set to the corresponding value in the input line.
+ . if the process didn't have an undo list, one will be created for it, and
+   updated according to the input line.
+
+Writing to this file is particularly useful if one wants to restore a process'
+state in terms of sysv ipc's, during a restart phase, for example.
+
+NOTE:
+The access flags used to open the semundo file are very important:
+opening the file in rd/wr mode will be considered as as "open for write".
+This means that a process opening another process' semundo file in rd/wr mode
+will see a failure, even if that process only wanted to read the file.
+

```

+2.18 /proc/sys/fs/types/ - File system type specific parameters

There's a separate directory /proc/sys/fs/types/<type>/ for each

--

Containers mailing list

Containers@lists.linux-foundation.org

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