
Subject: Re: [BUG][cryo] Create file on restart ?
Posted by [Matt Helsley](#) on Wed, 16 Jul 2008 20:59:23 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Wed, 2008-07-16 at 14:26 -0500, Serge E. Hallyn wrote:
> Quoting sukadev@us.ibm.com (sukadev@us.ibm.com):
> >
> > cryo does not (cannot ?) recreate files if the application created
>
> I think that's for the best.
>
> Don't you?

I agree. I think drawing the line for process checkpoint/restart before preserving the contents of mounted filesystems is very reasonable since mounted filesystem(s) can already be preserved with your choice of tool(s). I think it also gives us more options as far as using checkpointed images for error recovery; if we take the concept of checkpoint too far we may limit ourselves to merely reproducing errors rather than also giving ourselves a means to recover from errors.

Cheers,
-Matt Helsley

> -serge
>
> > a file before checkpoint and the file does not exist at the time
> > of restart.
> >
> > Note that the 'flags' field in '/proc/\$pid/fdinfo/\$fd' will not
> > have the O_CREAT (or O_TRUNC, O_EXCL, O_NOCTTY) flags. These
> > are cleared in __dentry_open()).
> >
> > At the time of restart, is there a way for cryo to know that the
> > file must be created ?
> >
> > To reproduce:
> > - run following program,
> > - checkpoint after the first printf
> > - rm /tmp/foo1
> > - restart # fails to open file during restart
> >
> > ---
> > #include <stdio.h>
> > #include <unistd.h>
> > #include <errno.h>
> > #include <sys/fcntl.h>
> >

```
> > main()
> > {
> > int fd;
> > int i;
> > char *buf = "abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ";
> >
> > fd = open("/tmp/foo1", O_RDWR|O_CREAT|O_TRUNC, 0666);
> >
> > if (fd < 0) {
> >     perror("open");
> >     exit(1);
> > }
> > printf("%d: Opened '/tmp/foo1', fd %d\n", getpid(), fd);
> >
> > for (i = 0; i < strlen(buf); i++) {
> >     if (write(fd, &buf[i], 1) < 0) {
> >         printf("Error %d writing %c to file, i %d\n",
> >             errno, buf[i], i);
> >         exit(1);
> >     }
> >     printf("%d: i %d, wrote %c\n", getpid(), i, buf[i]);
> >     sleep(2);
> > }
> > }
```

```
> > Containers mailing list
> > Containers@lists.linux-foundation.org
> > https://lists.linux-foundation.org/mailman/listinfo/containers
>
> Containers mailing list
> Containers@lists.linux-foundation.org
> https://lists.linux-foundation.org/mailman/listinfo/containers
```

```
Containers mailing list
Containers@lists.linux-foundation.org
https://lists.linux-foundation.org/mailman/listinfo/containers
```
