

Monday, September 1, 2008

Less known Solaris features: Getting rid of Zombies

Once in a while you will see some strange processes with instead of a process name. This happens when a child process terminates, but the parent process isn't interested in the outcome because it didn't wait for the child's response. Almost all resources of the child process are freed up at the moment with the exception of the entry in the process table. The parent process needs it to get the exit code from its child, thus you can't simply delete it on the termination of the child. The remaining process table entry will be deleted, when the parent process reaps the child process by gathering the exit code. But when the parent forgets to reap the child, it's undead, it's defunct. Or to stay in the terminology: You've produced a Zombie process.

Let's create such a process. It's really easy, we just have to create a long running process forking away a child but we don't use the wait() system call to gather its response at the exit.

```
bash-3.2$ nohup perl -e "if (fork())>0) {while (1) {sleep 100*100;}};"&Okay, let's check for our processes. In the output of ps -ecl the zombie processes are marked with a Z:
```

```
bash-3.2$ ps -ecl |grep "Z"
```

```
F S UID PID PPID CLS PRI ADDR SZ WCHAN TTY TIME CMD
0 Z 100 27841 27840 - 0 - 0 -? 0:00
```

bash-3.2\$A kill -9 to this process is without effect. Obviously, a zombie will go away when you terminate the parent process, but that isn't always an option. How can you get rid of this Zombies? Okay, with Solaris you can reap such processes manually. The preap forces the parent to reap the child by calling wait() system call on the child.

```
bash-3.2$ preap 27841
```

```
27841: exited with status 0And when you look in the process table again you will see, that the zombie found its peace...
```

```
bash-3.2$ ps -ecl |grep "Z"
```

```
F S UID PID PPID CLS PRI ADDR SZ WCHAN TTY TIME CMDObviously, you should ask yourself, why an application leaves such zombie processes, when the task of reaping them away manually gets a frequent task. Often it's because of bad programming style.
```

Posted by Joerg Moellenkamp in English, Solaris at 19:56

I knew there was a way I'd done this before, but for the life of me I couldn't remember how. I'd completely forgotten about the preap. Now to try and lodge it in my memory for next time.

Anonymous on Sep 2 2008, 16:36

Is there a way to do this prior to Solaris 9?

Anonymous on Sep 2 2008, 22:59

No, but then I hadn't thought of "preap".

In any case, other than to make your "ps" output cleaner, there is [almost] no need to reap these processes. The only thing that they occupy is a process table slot which is just a few bytes of memory.

It is a rare and troubled system that does not have a few (or a few thousand) process slots to spare.

Anonymous on Sep 3 2008, 22:32

Also, if you have a bunch of zombies and you want to kill them recursively, here's how I do it:

```
ps -A | grep defunct | awk '{print $1}' | xargs preap
```

Hope it helps. Cheers.

Anonymous on Nov 16 2012, 04:13

Or do something like ...

```
for i in $(ps -ecl | nawk '$2 == "Z" {print $4}'); do preap $i ; done
```

Anonymous on Apr 1 2014, 15:46

Funny thing, I just found out that there is no equivalent command for Linux.

Anonymous on Oct 13 2016, 15:51