

Friday, May 9, 2008

Less known Solaris features: About crashes and cores - Part 2: Forcing a dump

Okay, a dumps are not only a consequence of errors. You can force the generation of both kinds. This is really useful when you want to freeze the current state of the system or an application for further examination.

Forcing a core dumpLet's assume you want to have an core dump of a process running on your system:# ps -ef | grep "bash" | grep "jmoekamp"

```
jmoekamp 681 675 0 20:59:39 pts/1 0:00 bash
Okay, now we can trigger the core dump by using the process id of the process.# gcore 681
```

gcore: core.681 dumpedOkay, but the kicker is the fact, that the process still runs afterwards. So you can get an core dump of your process for analysis without interrupting it.# ps -ef | grep "bash" | grep "jmoekamp"

```
jmoekamp 681 675 0 20:59:39 pts/1 0:00 bash
Neat isn't it. Now you can use the mdb to analyse it, for example to print out the backtrace:# mdb core.681
```

```
Loading modules: [ libc.so.1 ld.so.1 ]
```

```
> $c
```

```
libc.so.1`__waitid+0x15(0, 2a9, 8047ca0, 83)
```

```
libc.so.1`waitpid+0x63(2a9, 8047d4c, 80)
```

```
waitjob+0x51(8077098)
```

```
postjob+0xcd(2a9, 1)
```

```
execute+0x77d(80771c4, 0, 0)
```

```
exfile+0x170(0)
```

```
main+0x4d2(1, 8047e48, 8047e50)
```

```
_start+0x7a(1, 8047eec, 0, 8047ef0, 8047efe, 8047f0f)
Forcing a crash dumpOkay, you can force a crash dump, too. It's quite easy. You can trigger it with the uadmin command.bash-3.2# uadmin 5 0
```

```
panic[cpu0]/thread=db47700: forced crash dump initiated at user request
```

```
d50a2f4c genunix:kadmin+10c (5, 0, 0, db325400)
```

```
d50a2f84 genunix:uadmin+8e (5, 0, 0, d50a2fac, )
```

```
syncing file systems... 2 1 done
```

```
dumping to /dev/dsk/c0d0s1, offset 108593152, content kernel
```

```
100% done: 31255 pages dumped, compression ratio 5.31, dump succeeded
```

```
Press any key to reboot.
```

Why should you do something like that? Well, there are several reasons. For example, when you want to stop a system right at this moment. There is an effect in clusters called "split brain". This happens, when both nodes of a cluster believe they are the surviving one, because they've lost the cluster interconnect(simplification warning). Sun Cluster can prevent this situation by something called quorum. In a high availability situation the nodes of a cluster try to get this quorum. Whoever gets the quorum, runs the service. But you have to ensure that the other nodes don't even try to write something to disks. The simplest method: Panic the machine.

Another use case would be the detection of an security breach. Let's assume, your developer integrated a security hole as large as the Rhine into a web applicaiton by accident and now someone else owns your machine. The false reaction would be: Switch the system off or trigger a normal reboot. Both would lead to the loss of the memory content and perhaps the hacker had integrated a tool in the shutdown procedure to erase logs. A more feasible possibility: Trigger a crash dump. You keep the content of the memory and you can analyse it for traces to the attacker.

Posted by Joerg Moellenkamp in Solaris at 13:01

BTW: pgrep and ps both take the -u argument, so you could write

```
pgrep -l -u jmoekamp bash
```

It might be less interesting for interactive process listing (limited columns) but good for scripting.

Anonymous on May 11 2008, 22:29

I don't use the pgrep command for some strange reasons ... my internal brain macros prefer the pipes ... obviously you are correct.

Anonymous on May 11 2008, 23:26