

Sunday, March 16. 2008

Less known Solaris Features: Resource Management - Part 3: Limiting operating system resources

The kernel of an operating system provides a huge amount of resources to the processes running on it. Such resources are file descriptors, shared memory segments or the process tables. Albeit it's hard to fill up this resources with modern operating systems it's not impossible. When a resource is consumed by a single malicious or erroneous application, all other can't run as well when they need more resources from the operating system.

Let's assume this scenario: There is an course "Perl Scripting for beginners" at the Unseen University and in the last year the lesson "About fork" ended in chaos as some of your students coded forkbombs as they though this would be funny (as the class before, and the one before ...)

```
#!/usr/bin/perl
```

fork while 11've stored this little script at /opt/bombs/forkbomb.pl. A few seconds after starting such a script, the system is toast because of the hundreds of forked processes. Don't try this without resource management. Okay, but this year, you've migrated to Solaris. You can impose resource management.

Okay, we have to modify our project configuration:

```
# projmod -K "task.max-lwps=(privileged,10,deny)" class2005
```

Now we have configured a resource limit. A single task in the class2005 cant have more than 9 processes. The tenth attempt to fork will be denied. Okay, do you remember the reasons, why the system starts a new task? One of it is "login". Thus every login of a user gives him 10 threads to work with. And this is exactly the behaviour we want.

Let's assume Alice starts her forkbomb: # ps -ef | grep "alice"

```
alice 685 682 0 14:58:12 ? 0:00 /usr/lib/ssh/sshd
alice 693 686 14 14:58:42 pts/1 0:38 /usr/bin/perl /opt/bombs/forkbomb.pl
alice 686 685 0 14:58:12 pts/1 0:00 -sh
alice 694 693 15 14:58:42 pts/1 0:38 /usr/bin/perl /opt/bombs/forkbomb.pl
alice 695 694 14 14:58:42 pts/1 0:37 /usr/bin/perl /opt/bombs/forkbomb.pl
alice 696 695 14 14:58:42 pts/1 0:37 /usr/bin/perl /opt/bombs/forkbomb.pl
alice 697 696 14 14:58:42 pts/1 0:37 /usr/bin/perl /opt/bombs/forkbomb.pl
alice 698 697 14 14:58:42 pts/1 0:39 /usr/bin/perl /opt/bombs/forkbomb.pl
alice 699 698 14 14:58:42 pts/1 0:38 /usr/bin/perl /opt/bombs/forkbomb.pl
```

```
# ps -ef | grep "alice" | wc -l
```

9After forking away 7 forkbomb.pl processes, any futher fork is denied by the system. The load of the system goes up (as there are hundreds of denied forks) but the system stays usable.

Alice sends her script to Bob. He tries it, too:# ps -ef | grep "alice"

```
alice 685 682 0 14:58:12 ? 0:00 /usr/lib/ssh/sshd
alice 28520 28519 6 15:15:08 pts/1 0:03 /usr/bin/perl /opt/bombs/forkbomb.pl
alice 686 685 0 14:58:12 pts/1 0:00 -sh
alice 28521 28520 6 15:15:08 pts/1 0:03 /usr/bin/perl /opt/bombs/forkbomb.pl
alice 28519 686 6 15:15:08 pts/1 0:02 /usr/bin/perl /opt/bombs/forkbomb.pl
alice 28522 28521 6 15:15:08 pts/1 0:03 /usr/bin/perl /opt/bombs/forkbomb.pl
alice 28524 28523 6 15:15:08 pts/1 0:03 /usr/bin/perl /opt/bombs/forkbomb.pl
alice 28523 28522 6 15:15:08 pts/1 0:03 /usr/bin/perl /opt/bombs/forkbomb.pl
alice 28525 28524 6 15:15:08 pts/1 0:02 /usr/bin/perl /opt/bombs/forkbomb.pl
```

```
# ps -ef | grep "bob"
```

```
bob 28514 28511 0 15:14:47 ? 0:00 /usr/lib/ssh/sshd
bob 28515 28514 0 15:14:47 pts/3 0:00 -sh
bob 2789 2502 6 15:15:10 pts/3 0:03 /usr/bin/perl /opt/bombs/forkbomb.pl
bob 2791 2790 6 15:15:10 pts/3 0:03 /usr/bin/perl /opt/bombs/forkbomb.pl
bob 2502 28515 6 15:15:10 pts/3 0:03 /usr/bin/perl /opt/bombs/forkbomb.pl
bob 2790 2789 6 15:15:10 pts/3 0:03 /usr/bin/perl /opt/bombs/forkbomb.pl
bob 2792 2791 6 15:15:10 pts/3 0:03 /usr/bin/perl /opt/bombs/forkbomb.pl
bob 2793 2792 6 15:15:10 pts/3 0:02 /usr/bin/perl /opt/bombs/forkbomb.pl
bob 2794 2793 6 15:15:10 pts/3 0:03 /usr/bin/perl /opt/bombs/forkbomb.pl
```

This is still no problem for the system. After a few forks of the forkbomb, the system denies further forks. And the system stays usable. The limitation

of the number of processes is only one example. You can limit other resources. You can find a list of all controls at the man page of `resource_controls`

Posted by Joerg Moellenkamp in English, Solaris at 22:33