
Subject: Re: Debian-style init scripts considered harmful?

Posted by [Steve Wray](#) on Thu, 17 Jul 2008 01:41:24 GMT

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

Kir Kolyshkin wrote:

> Steve Wray wrote:

>> Hi there,

>>

>> Debian uses start-stop-daemon in the init scripts to, for one thing,

>> stop services.

>>

>> From the man page:

>>

>> Note: unless --pidfile is specified, start-stop-daemon behaves similar

>> to killall(1). start-stop-daemon will scan the process table looking

>> for any processes which match the process name, uid, and/or gid (if

>> specified). Any matching process will prevent --start from starting the

>> daemon. All matching processes will be sent the KILL signal if --stop

>> is specified. For daemons which have long-lived children which need to

>> live through a --stop you must specify a pidfile.

>>

>> For example, nfs-kernel-server does not use --pidfile. It looks for

>> nfsd processes to kill.

>>

>> Suppose that the Openvz host and one of its guests were running NFS

>> and, on the host, one were to run /etc/init.d/nfs-kernel-server stop

>>

>> As I understand it this would have the side-effect of killing off the

>> nfsd processes on the guest.

>>

> That is right, and this is just one of the reasons why we don't

> recommend to run anything (but the needed bare minimum like sshd) on the

> host system.

In my case, this isn't practical; I use cfengine to manage and maintain
virtually all of our servers. We have a lot of servers.

In fact, it was cfengine which brought this to my attention; I restarted
it on the openvz host and then started to get nagios alerts about
cfengine not running on any of the guests.

It was at this point that I realised that openvz isn't a virtualisation
environment; its a very *very* sophisticated chroot.

> There is a solution and a workaround for the problem. The solution is,
> right, to fix bad initscripts. I mean, it's not OpenVZ-specific --
> relying on process names is wrong, any user can run a process named nfsd

> and it should not be killed.

>

> The workaround is to introduce a feature to hide guests' processes from
> the host system. This is implemented in OpenVZ kernels $\geq 2.6.24$ as per
> bug #511 (<http://bugzilla.openvz.org/511>).

Well I look forward to trying this out some time!
