

setting the core dump name schema

Why would you want to change the names of the coredump from the default?

Update: this is also a very useful defence against a currently known severe security affecting bug. See Core Dump Handling Vulnerability for information and workarounds

Lets assume you got a couple of core dumps files in one directory (e.g. home) and you don't know which program caused it.

Unless you have a grsec/pax kernel with logging features enabled, it might get tricky to find out who caused the coredumps.

This is the point where the naming scheme comes handy :)

The pattern can be read/set via `/proc/sys/kernel/core_pattern`.

To set a new one just do:

```
echo "newpattern" > /proc/sys/kernel/core_pattern
```

you can have variables to make the file named different per executable, pid a.s.o..
max length 64 characters; default value is "core"

Here is a small list of possible variables

```
%p:      pid
%>NUL>:  '%' is dropped
%%:      output one '%'
%u:      uid
%g:      gid
%s:      signal number
%t:      UNIX time of dump
%h:      hostname
%e:      executable filename
%>OTHER>: both are dropped
```

If `core_pattern` does not include "%p" (default does not) and `core_uses_pid` is set, then .PID will be appended to the filename.

example:

```
echo "core.%e.%p" > /proc/sys/kernel/core_pattern
```

produces files names

```
core.>executable>.>pid>
```

to make the changes permanent, add the following line to `/etc/sysctl.conf`:

```
kernel.core_pattern = core.%e.%p
```

happy core dumping.

Support this site