CGroups: Subsystems as Modules

bblum (Ben Blum)

What is this 'cgroups'?

- Part of the core linux kernel
- Group, manage, limit processes
  - cpuacct, cpuset, memory, nsproxy, devices
  - “Subsystems”
- kernel/cgroup.c, include/linux/cgroup.h, ...

cgroups interface

```
# cd /dev; mkdir cgroup
# mount -t cgroup -o SS1,SS2,... \
  none cgroup/
# ls cgroup/
  subsys1.*  tasks
  subsys2.*  procs
  subsys3.*
```

Goal (why?)

- Make subsystems compilable as modules
- Why?
  - Dynamic loading/unloading as needed
  - Build subsystems outside of the kernel

Goal (how? (vague))

- Make subsystems compilable as modules
- How?
  - Within cgroups:
    - cgroup_init_subsys()
    - struct cgroup_subsys *subsys[]
    - for_each_subsys()
  - Module-wise:
    - initcall() (for module init/exit)
  - Kernel module build system (menuconfig, ...)

- http://cgrouphacking.blogspot.com/