Process Management under Linux

Mir Farooq Ali

Process States

- Recall:
  - A parent process can suspend a child process
- Therefore, if a child is in run state and goes to ready (time slice up), and the parent runs and decides to suspend the child, then how do we reflect this in the process state diagram ???
- We need 2 more states
  - Ready suspended
  - Blocked suspended

Process State diagram reflecting Control

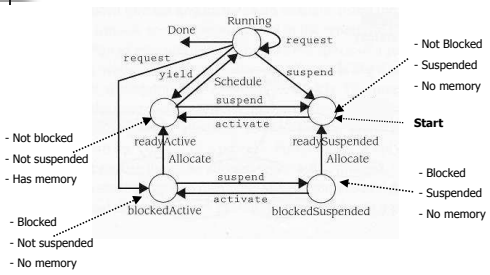

Processes in Linux

- Also called *tasks*
- Task table or process table defined in src/linux/include/sched.h

```
extern struct task_struct
  *pidhash[PIDHASH_SZ];
```

- Can also be accessed as a doubly-linked list `p->next_task` and `p->prev_task`

Process or task descriptor

- Called task_struct
- Present in src/include/linux/sched.h
- Contains various fields to indicate
  - state
  - priority
  - pointers to parent, children, other tasks in pid list
  - tty
  - memory location
  - file descriptors
  - …

Process States

- Linux identifies following states

1. `TASK_RUNNING`
2. `TASK_INTERRUPTIBLE`
3. `TASK_UNINTERRUPTIBLE`
4. `TASK_ZOMBIE`
5. `TASK_STOPPED`
6. `TASK_EXCLUSIVE`

Process Creation

- Remember in traditional UNIX, we use fork() and then typically exec()
- fork() duplicates resources owned by parent for child process and copies them to new address space
- This method is slow and inefficient, since exec() wipes out address space anyway

Process creation in Linux

- Copy On Write technique
- Lightweight processes
- vfork()

Copy-on-write

- Child pages are pointers to parent pages
- If child makes a change to a page, a new copy is made for the child
- This way, you avoid making separate copies of pages unnecessarily

Lightweight processes

- Allow parent and child processes to share many kernel data structures
- created in Linux by function called __clone()
- uses non-standard clone() system call

vfork()

- Creates a process that shares memory address of parent
- Parent is blocked until child exits or executes a new program by doing exec()

User view of processes

- Can use ps command with various options, for example,
  - ps –aux
  - ps –ef

/proc file system

- process information pseudo file system
- Do `man proc` to get more info
- /proc directory contains
  - Numerical subdirectory for each running process
  - A number of other files containing kernel table information

/proc… continued

- Files include
  - cpuinfo – contains CPU specs
  - uptime – time in secs since machine was last rebooted and idle time since then
  - version – kernel version
  - loadavg – Load average of machine over the past 1, 5 and 15 minutes
  - …

Process directories

- One subdirectory for each running process
- Files include
  - cmdline
  - cwd
  - environ
  - exe
  - fdm
  - map
  - mem
  - root

References

- Linux Kernel 2.4 internals, Tigran Aivazian http://www.tldp.org/LDP/lki/
- Modern Operating Systems, 2nd Ed., A. Tanenbaum
- Understanding the Linux Kernel, D. Bovet, and M. Cesati