




DREAMWORKS™
— SKG —

About Box

- I'm just a SysAdmin
 - 23 years experience with Unix/Linux
 - I've hacked kernels
 - but I'm not a kernel hacker
 - 1+ years with PDI/Dreamworks
- PDI/Dreamworks Animation
 - We make animated movies
 - Shrek, Madagascar, Flushed Away, Over the Hedge...
- We have two sites
 - Southern California (2000+ Farm, 800+ Desktop, 200+ TB)
 - Northern California (700+ Farm, 400+ Desktop, 100+ TB)



About (con't)

- Common Compute Environment
 - Mostly multi-core, multi-processor, 64-bit
 - 4 gig desktops, 8 gig render farm machines
 - Extreme reliance on NFS and the Automounter
- Process Mix
 - Both batch and interactive processes are very large
 - We run into swap **on most machines daily**
 - We run out of memory often
 - Render batch jobs on desktops at night
 - Batch jobs push interactive process to swap
 - OOM: Batch jobs continue, interactive processes get killed

Memory Related Issues

- Control which processes are terminated during memory starvation
 - Not just to kill or not kill, or system lockup
 - Not a compile time option (OOM_DISABLE is good!)
- No way to quantify process memory usage
 - Needs to be lightweight (not a kernel debugger)
 - Prefer not to have to count pages ourselves
 - Needs to handle shared/mapped/swapped pages
- We have no way to contain a set of processes to a maximum memory footprint
 - RLIMIT_DATA didn't do what we thought
- Overnight batch jobs shove interactive processes into swap, causing user unhappiness at swap-in time

Other Issues

- NFS Mounting
 - We mount the same Filesystems in different locations on clients, with different permissions
 - Behavior changed between releases of 2.6 kernel
 - Can no longer do this, thus breaking some security models.