



CE Linux Forum

Linux-tiny revival

Tim Bird

Chair, Architecture Group
of the CE Linux Forum

(for Japan Technical Jamboree #16, Aug 31, 2007, Tokyo)



Linux-tiny revival

- Introduction to Linux-tiny project
- History
- Status
- Recent work



Intro to Linux-tiny

- What is Linux-tiny?
 - A collection of patches to reduce the memory size of the Linux kernel
 - Currently, a collection of about 40 patches
 - Had as many as 110 patches at one time
 - Some patches mainlined, some are out-of-date



Linux-tiny Patch Areas

- Configurable static data structure size
- Configurable code elimination
 - Eliminate a subsystem or feature, depending on kernel configuration
- Moving inlines into stand-alone functions
- Kernel instrumentation and analysis tools for memory usage reporting



Linux-tiny History

- Started in 2003 by Matt Mackall
- CELF sponsored Matt in 2005/2006
- Matt mainlined the top 17 patches from Linux-tiny
- Project was mostly abandoned in 2006
 - For a while, was not maintained relative to recent kernels



Current status

- CELF is reviving the project
- New maintainer volunteered at Ottawa Linux Symposium, 2007:
 - Michael Opdenacker
- Michael is an experienced embedded Linux trainer
 - Has his own training company, Free Electrons



Status (cont.)

- Have not formally announced the new maintainer
 - Watch celinux-dev or kernel mailing list next week
- Lots of work recently by Sony
 - Up-port to 2.6.22.5
 - Converted kmallocc-accounting feature to slab-accounting
 - Fixed lots of bugs
 - instrument all cache-related dynamic allocations



Results

Config option	PC size	OSK size
NET_SMALL	8192	6368
PROC_KCORE	2556	0 ?
SYSEENTER	7029	not applicable
AIO	8114	10396
XATTR	3780	0 error (have fix)
FILE_LOCKING	116228	442840 (error in test code?)
ETHTOOL	5340	7328
INETPEER	2848	3188
NET_SK_FILTER	3988	4008
NET_DEV_MCAST	2098	2152
IGMP	0 error	13112
BINFMT_ELF_AOUT	1416	2032
IDE_HWIFS	9216	6208
(others: MAX_SWAPFILES_SHIFT, NR_LDISCS, MAX_USER_RT_PRIO,SERIAL_PCI, PCI_QUIRKS)	25456	5152
Total	196261	59944



Slab Accounting Feature

- Records all callers to dynamic memory allocation system (slab cache) in kernel
- Tracks statistics for each caller
- To use:
 - Turn on `DEBUG_SLAB` and `DEBUG_SLAB_ACCOUNT`
 - On the Kernel Hacking menu
 - `cat /proc/slab_account`



Slab Info Analysis Programs

- Slabalyze
 - Collects information from `/proc/slabinfo`
 - Shows overhead used by different caches
 - Can sort by total memory, overhead, etc.
- casort - Cache accounting sort
 - Collects information from `/proc/slabaccount`
 - Shows callers, sorted by different fields



Demonstration

- On OSK machine in US test lab
 - if we are lucky...



What is next?

- Update documentation:
 - Linux kernel size tuning guide
 - Linux-tiny patch status matrix
 - Have automated tool for testing Linux-tiny patches
- Want to get `cache_accounting` working on slob allocator
- Want runtime testing of existing patches
- Please download and test the code
 - Every bug report is important
- Will start trying to mainline sub-patches again, this fall!



Resources

- Web site:
 - http://elinux.org/Linux_Tiny
- Instructions:
 - http://elinux.org/Kernel_Size_Tuning_Guide
 - Will update soon (goal is by next week)
- Mailing List:
 - linux-tiny@selenic.com



Open Discussion