



Android OS for Servers?

John Stultz
April 13th 2011

<http://www.linaro.org>



Android™ OS for what?



Portions of this page are modifications based on work created and shared by Google and used according to terms described in the Creative Commons 3.0 Attribution License.
See: <http://www.android.com/branding.html>

<http://www.openclipart.org/detail/mainframe-by-kattekrab>

Android™ OS for what?



+



Portions of this page are based on work created and shared by Google and used according to terms described in the Creative Commons 3.0 Attribution License.
See: <http://www.android.com/branding.html>

<http://www.openclipart.org/detail/mainframe-by-kattekrab>



What is the Android OS?

- Linux™ kernel
- Android patches
- Bionic libc
- Other libraries
- Dalvik vm
- Application frameworks
- Apps



What is the Android OS?

- Linux™ kernel
- Android patches
- Bionic libc
- Other libraries
- Dalvik vm
- Application frameworks
- Apps



What's in the Android Patches?

- Ashmem
- Binder
- Pmem
- Logger
- Early suspend
- Wakelocks
- Alarm Timer
- LowMemoryKiller
- Paranoid network
- Yaffs2 fs
- Ram_console
- Apanic
- Adb gadget driver
- Gpio patches
- Lots of other small fixes and hacks for arm, mmc, Bluetooth™, etc.




What's in the Android Patches?

- ~249 patches in total for the 2.6.38 tree
 - 3.3 megs of patches
- Each patch (well, mostly) represents a limitation that the Android developers found in the Linux kernel



Goals of Android Patches:

- Fix bugs, enable new hardware
- Improve power management
- Improve error reporting
- Increase security
- Improve performance



Fixing bugs/Enable new hardware

- Ashmem
- Binder
- Pmem
- Logger
- Early suspend
- Wakelocks
- Alarm Timer
- LowMemoryKiller
- Paranoid network
- Yaffs2 fs
- Ram_console
- Apanic
- Adb gadget driver
- Gpio drivers
- Lots of other small fixes and hacks for arm, mmc, Bluetooth, etc.



Power management

- Ashmem
- Binder
- Pmem
- Logger
- **Early suspend**
- **Wakelocks**
- **Alarm Timer**
- LowMemoryKiller
- Paranoid network
- Yaffs2 fs
- Ram_console
- Apanic
- Adb gadget driver
- Gpio patches
- Lots of other small fixes and hacks for arm, mmc, Bluetooth, etc.



Error reporting

- Ashmem
- Binder
- Pmem
- **Logger**
- Early suspend
- Wakelocks
- Alarm Timer
- LowMemoryKiller
- Paranoid network
- Yaffs2 fs
- **Ram_console**
- **Apanic**
- Adb gadget driver
- Gpio patches
- Lots of other small fixes and hacks for arm, mmc, Bluetooth, etc.



Security

- Ashmem
- Binder
- Pmem
- Logger
- Early suspend
- Wakelocks
- Alarm Timer
- LowMemoryKiller
- Paranoid network
- Yaffs2 fs
- Ram_console
- Apanic
- Adb gadget driver
- Gpio patches
- Lots of other small fixes and hacks for arm, mmc, Bluetooth, etc.



Improve performance

- Ashmem
- Binder
- Pmem
- Logger
- Early suspend
- Wakelocks
- Alarm Timer
- LowMemoryKiller
- Paranoid network
- Yaffs2 fs
- Ram_console
- Apanic
- Adb gadget driver
- Gpio patches
- Lots of other small fixes and hacks for arm, mmc, Bluetooth, etc.




When will it all be merged?

- It is already going upstream
 - 100 commits with @android.com authors already in mainline *
 - Over half of those are from 2011 already
 - Mostly ARM support
- However, some items have met quite a bit of resistance



Talked about at length on lkml



*Includes variants like wake_lock and alarmtimer



Background on wakelocks

- Suspend mode powers down all hardware except memory
- One approach for optimized power management would be to suspend whenever idle
- Wakelocks are the primitive used to determine if the system is in use



Hindsight 20/20

- Wakelocks are a lot of change to solve a problem isn't thought of as a major issue with the upstream kernel
- Mainline community mostly focused on runtime power optimization
- Misunderstanding of both the problem and solution devolved into a take-it-or-leave-it situation, with sour feelings on both sides



Even so...

- Community is still moving forward
- pm_stay_awake/pm_relax kernel infrastructure has been merged
 - Similar to wakelocks
 - Unclear if sufficient for Android
- Android developers have even hinted that they may drop early suspend in the future



So back to servers...

- How do I suspend my system to save power, and still have my nightly backups run?
 - `/dev/rtc/rtcN` allows for this, but its messy
 - Just want a timer that fires regardless if the system is suspended or not, which then kicks off a chron job
- Need: Easy to use, programmatic way to wake a system
- Android Alarm Timer provides this



Android Alarm Timers

- Driver that provides a chardev interface
- /dev/alarm
- Use ioctls to read time and set alarms
- Duplicates some of the timekeeping interface, using new names and terminology



Posix Alarm Timers

- Different from Android Alarm Timers, in that it uses the Posix clock/timers API instead of `/dev/alarm` ioctl interface
 - `clock_nanosleep(CLOCK_REALTIME_ALARM,...)`
- Use the RTC hardware to allow for timers that fire regardless of if the system is suspended or not
- Allows for multiple events to be multiplexed upon a single RTC device



Posix Alarm Timers

- While the user-level interface is different than the Android Alarm Timers, the back end will be able to be shared
- Still a few open questions to decide:
 - What capabilities will be needed to set an Alarm timer?
 - What policy will there be to control alarm timers? (ie: avoid laptop overheating while in a bag)



Posix Alarm Timers

- Regardless of interface details, Alarm Timer functionality can allow servers and desktops to go to sleep and save electricity, while still allowing work scheduled in the future to run
- Could combine with wake-on-directed-packet functionality to allow web/interactive servers to suspend when idle



Posix Alarm Timers

- More interesting future question:
 - Once we can programmatically wake a system, can we programmatically make it sleep?
 - How would we know if was safe to suspend the system?
 - Might need some way to inhibit suspend...



LowMemoryKiller

- As opposed to the Out-Of-Memory (OOM) killer
- Kill applications before system observes negative effects prior to OOM state
- Obviously useful for servers
 - Web service vs. low-priority background jobs
 - Multiple VM guests with different priorities



Similar works in progress

- Ashmem
- Binder
- Pmem
- Logger
- Early suspend
- Wakelocks
- Alarm Timers
- LowMemoryKiller
- Paranoid network
- Yaffs2 fs
- Ram_console
- Apanic
- Adb gadget driver
- Gpio patches
- Lots of other small fixes and hacks for arm, mmc, Bluetooth, etc.



Other efforts



What about those fixes?

- Ashmem
- Binder
- Pmem
- Logger
- Early suspend
- Wakelocks
- Alarm Timer
- LowMemoryKiller
- Paranoid network
- Yaffs2 fs
- Ram_console
- Apanic
- Adb gadget driver
- Gpio patches
- Lots of other small fixes and hacks for arm, mmc, Bluetooth, etc.



Trivial tree

- Try to address the concern that good fixes in the Android tree were being overlooked due to controversy surrounding other unrelated Android patches
- Picking small changes that do not introduce major functionality
- Broken up into branches by “topics”
 - ARM, sched, mmc, Bluetooth, etc.
- Sending small patch queues to topic maintainers/lists for review



Trivial tree difficulties

- Not much uptake so far. Most patches considered by maintainers as short-term fixes, rather than comprehensive solutions
- Not being the patch author, its difficult to advocate for random patches in random subsystems
 - Can only try to facilitate the discussion
- Slow going, as keeping all the topic branches current is difficult



Linaro™ Android Platform

- Upstream focused Android community for ARM
- Linaro toolchain packages for Android
- Sophisticated infrastructure with cloud based build service and ARM board farm for automated testing, validation and benchmarking
- Melting pot for upstreaming of member SoC platform code



Linaro™ Android Platform

- Installable images and public code for low-cost Linaro member SoC boards
- Combining latest publicly available AOSP platform code with latest Linaro kernel and toolchain along with member SoC platform code
- Releasing monthly
- If you are interested in helping, please come talk to us in [#linaro-android](#) on freenode!



Summary

Development goals on embedded and server are really not so different.



Legal

This work represents the view of the authors and does not necessarily represent the view of IBM.

IBM is a registered trademark of International Business Machines Corporation in the United States and/or other countries.

Android is a trademark of Google Inc. Use of this trademark is subject to Google Permissions.

Other company, product, and service names may be trademarks or service marks of others.



Questions?