Cryptographic Hardware Support for the Linux Kernel

James Morris
Red Hat Inc.

Oregon Networking Summit, July 2004

Current Status

- Simple crypto API in the 2.6 kernel, designed primarily for IPSec and then disk encryption.
- Does not support hardware crypto at all.
- Optimized ASM modules should be supported soon in a simple way (i.e. automatically selected).
- Good range of algorithms, more than enough for IPSec and other users so far. Easy to add new algorithms, basically working OK.

Summary of future directions

- Optimized ASM modules. Simple support should be available soon where ASM module is automagically selected at kernel configuration time.
- Asymmetric crypto.
- Support for various hardware devices.
- Userspace API for access to hardware devices.

Optimized ASM

- Simple support is easily integrateable (this word should be shot) into the current 2.6 kernel.
- i586 AES ASM module from Fruhwirth Clemens is ready to be integrated, should be available soon.
- Config-time algorithm selection only for 2.6.
- As part of the hardware API, I would like to implement a more sophisticated runtime algorithm selection system, which tests the speed of all available implementations and tries to do the best thing by default. Admin override via API essential.

Asymmetric Crypto

- This is potentially useful in the kernel for verification of loadable modules and program loading.
- Some work happening in this area. Software support should be mergable for 2.6. Crypto API may not need to be changed very much.
- Asymmetric crypto is becoming more common in hardware, and is likely to offer very good benefits over software implementations compared to symmetric crypto.
- Userspace access to asymmetric crypto hardware would be useful, for e.g. SSL, SSH, IKE etc.

Hardware Devices

- Several types of cryptographic hardware:
  - PCI cards with crypto processors. Becoming faster and increasingly sophisticated.
  - Chipset/platform crypto (e.g. TPM, s390 z990).
  - CPU level, such as the VIA xcrypt instructions.
  - Devices with crypto processors integrated, such as the PRO100/S.
  - Level of sophistication ranges from simply performing AES or DES synchronously to programmable devices with asynchronous and parallel processing, and protocol offload.

Why Crypto Hardware? (1)

- Modern CPU performance for software crypto is very good, better than most crypto cards I have.
- Crypto processor cards will also require increased PCI bus trips.
- But, there are several cases for hardware crypto:
  - Scalability: high end system with good I/O may have several cards offloading cpu intensive tasks. Trade latency off against scalability.
  - Useful for specialized embedded systems running Linux.

Why Crypto Hardware? (2)

- Research paper by OpenBSD developers[1] has some very useful information (PCI based crypto processors):
  - Smaller buffers do not tend to benefit from hardware acceleration. Should be done in software or batched to hardware.
  - As buffers become larger, PCI bus transaction overheads are amortized. Dramatic performance improvement for larger buffers.
  - Aggregate performance very good with parallelization & multithreading (but higher latency).
  - Reduces CPU contention.

[1] http://www.openbsd.org/events.html#usenix2003

Requirements

- Asynchronous kernel crypto API:
  - Configurable batching, load balancing, parallel access (crypto scheduler?).
- Support for synchronous hardware (VIA).
- Driver API for crypto hardware.
- Configurable algorithm selection.
- User API for access to hardware (cryptoapifs?).
- Metrics.

Also see http://samba.org/~jamesm/crypto/hardware_notes.txt

What's needed.

- Someone to do the work, both design and coding. Not a trivial project.
- More hardware and documentation.
- GPL drivers.
- Mailing list to coordinate/discuss?

Current status of HW support

- Some high level requirements (per previous slide & web page).
- A few GPL drivers available, some incomplete. Much rework probably needed.
- Some hardware documentation.
- Several people have hardware, can probably get more.
- A company is reportedly working on this, but I have not heard anything from them in a while.
- Lots of people email me asking if anyone is working on it, because they want to. Nothing ever happens.
- Mailing list set up by Michael Ludvig, has five messages so far. http://lists.logix.cz/mailman/listinfo/cryptoapi

Issues / Discussion

- Protocol offload.
- Requirements capture important, please email them to me if you have any.