Nfnetlink story

Trying to improve nfnetlink_queue

Who is using it ?

- Known projects using libnetfilter_queue:
  - nufw
  - snort_inline
  - some IPS (notably SIP)

The problems

- Performance is the issue
  - Number of queued packets per second is far too low
  - 6000 pckts/s is badly reached

Trying to improve things

- Implementing buffered kernel to userspace messages:
  - less messages, less switch
  - should be better
- Nfnetlink_log has something similar:
  - queue threshold
  - timeout
- Solution is insufficient
  - delay is not admissible for all applications

Porting log to queue

- Files are really similar
  - Could it be factorized ?
- Port was simple
- Results were astonishing:
  - Reasonabily bad performance
  - Unable to scp through nfqnl_test

Going from queue to log

- libnfnetlink was buggy:
  - unable to handle packet containing multiple messages
- This was just the first step:
  - message type for last message in packet is overwritten to MSG_DONE.
  - libnfnetlink ignore message with this type.
  - with a queue threshold of 10 you only get 9 packets ...

Final result

- With buffered messages:
  - nfqnl_test resists to hping
  - performance are almost equal to standard version
  - in current version scp has some problems
- Seems to be useless:
  - no performance improvement
  - loss in transit time
- Last hope: buffering answer message