VoIP in Linux

Haifux

February 2010
Rami Rosen

ramirose@gmail.com

- Why VOIP ? Motivation.
  - Taking advantage of existing infrastructure of an organization.
    - LAN, Internet connection, domain names, desktops, etc.
  - Centralized management (mail when missing a call, fax to email, etc).
  - First VoIP deployments in Israel was around 2000.
    - “Given Imaging” and others.

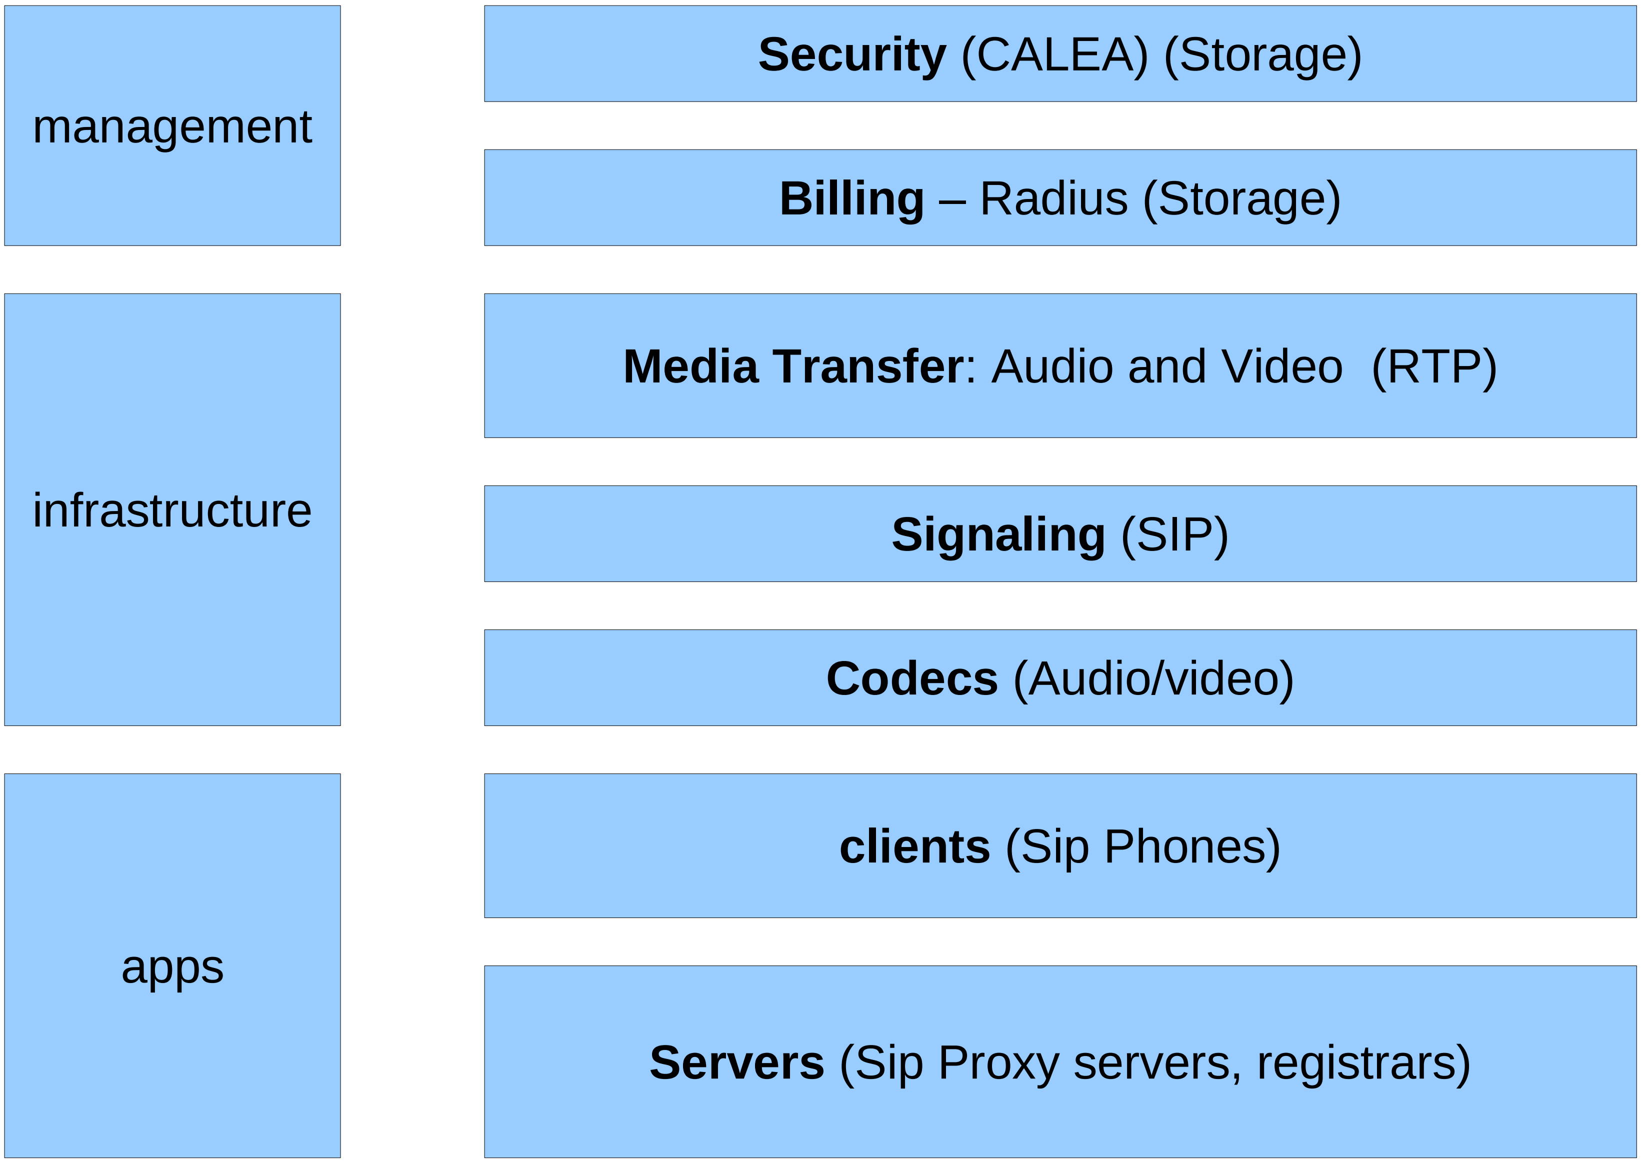
management
Security (CALEA) (Storage)
Billing – Radius (Storage)
infrastructure
Media Transfer: Audio and Video (RTP)
Signaling (SIP)
Codecs (Audio/video)
apps
clients (Sip Phones)
Servers (Sip Proxy servers, registrars)

- Most VOIP applications use the **RTP** protocol and the **SIP** protocol nowadays.
- **RTP – Real Time protocol**
- **SIP - Session Initiation Protocol**

RTP

- The Media flow is using RTP protocol.
  - RTP: Real-Time Transport Protocol.
    - Data flow (RTP packets)
    - Control flow (RTCP packets).
      - Ekiga client does work with RTCP.
      - Not all application implement RTCP.
    - Video conferencing uses two RTP sessions:
      - Send video on one RTP port
      - Send audio on a second RTP port.
        - Lip Synchronization between two sessions (using info from RTCP)

- Why two sessions ?
  - Enable separate QoS for each session
  - Enable the other side to hear only audio or view only video
  - More simple to implement

- Media in RTP is sent usually by UDP over IP.
  - However, UDP isn't mandatory according the spec does not.
  - So also sending media in RTP by TCP is acceptable according to the spec.

Diagram of media RTP session

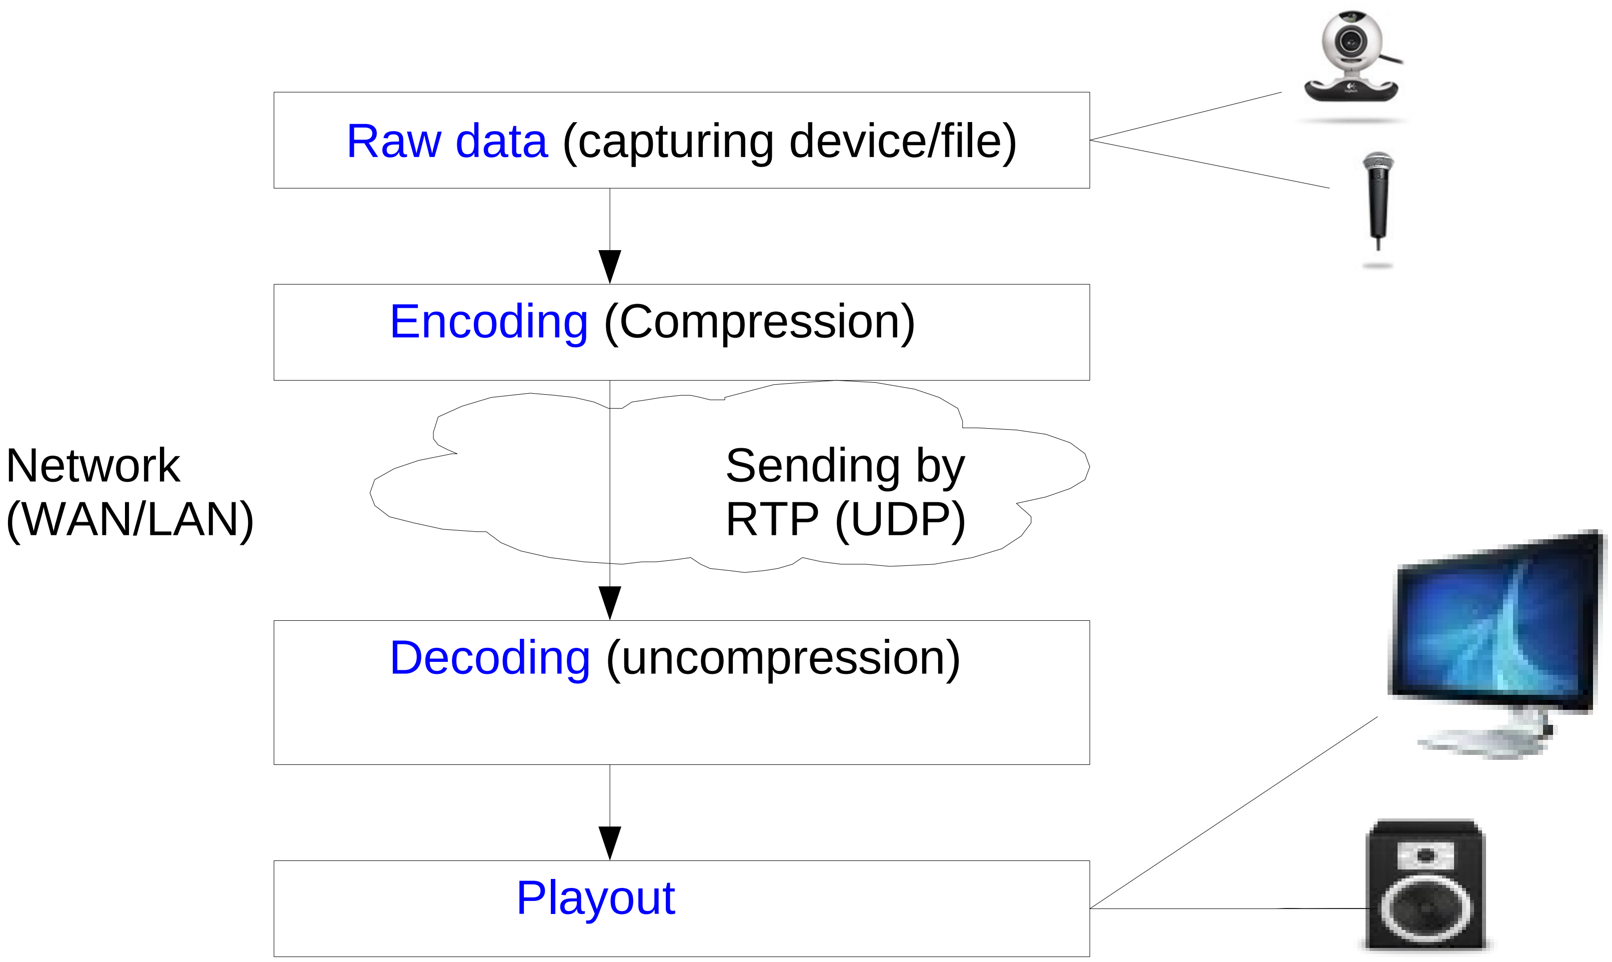

- The encoding is much more heavier in terms of cpu usage than decoding. Therefore there are cases when encoding is performed in hardware (Special chips, dsp, etc.)
  - Sometimes integrated within the microphone/webcam
- Codecs is use:
  - Audio: G.711, G.723, G.726, G.729, GSM, iLBC, AMR, speex. (There are a lot more).
  - Video: h.261, h.263, h.264.

- Many codecs support **VAD** (Voice Activity Detection)
  - This enables to stop sending RTP packets during silent periods or send little zeroed packets (with payload size of 4 bytes).

- Audio codecs:
  - Most basic is G.711
    - The input sample is 16 bit raw (PCM).
    - 8000 samples in a second (8Khz).
    - After encoding, each sample becomes 8 bits.
    - This means the the bitrate is **64 Kb/s.**

- G.729
  - The input sample is also 16 bit raw (PCM).
  - 8000 samples in a second (8Khz).
    - In one ms we have 8 samples
    - In 10 ms we have 80 samples.
    - 10 ms size (after encoding) is 10 bytes.
    - This means the bitrate is **8 Kbit/s.**
      - Which is 8 times better than 711 (for less quality).
      - Sometimes the frames sent are of 20 ms
        - 50 packets in a second.

- AMR
  - Adaptive Multi Rate
    - Bitrates 4.75, 5.15, 5.90, 6.70, 7.40, 7.95, 10.2, or 12.2 kbit/s
    - AMR codec is used also as a cellular phone code by many cellular companies.
      - For example, Pelephone in Israel uses AMR in its UMTS.
      - Some cellular companies intend to use also WB-AMR
        - WideBand AMR.
    - Error correction algorithms.

- The range of the human voice extends **from 80 hertz to 14,000** hertz.
- Wideband audio gives a higher quality voice.
- Wideband is also sometimes called "HD Voice"

RTP packet

<table>
<tr><td>V</td><td>P</td><td>X</td><td>CC</td><td>M</td><td>PT</td><td>Sequence Number</td></tr>
<tr><td colspan="7">TimeStamp (32 bit)</td></tr>
<tr><td colspan="7">Synchronization Source (SSRC) identifier (32 bit)</td></tr>
<tr><td colspan="7">Contributing source (CSRC) identifier (32 bit)</td></tr>
</table>

- V (2 bits) = version (value is always 2)
- P (1 bit) – padding.
- X (1 bit) - extension.
- CC (4 bits) - CSRC count
- M - 1 bit – marker.
  - Maker is usually 0. When there are special events, like DTMF, it is 1.
  - Also when the last fragment of a video packet which was fragmented.

- PT (7 bits) - Payload Type
  - Many apps use PT according to RFC 1890.
    - “RTP Profile for Audio and Video Conferences with Minimal Control”. (RFC 1890.)
    - For example, GSM has PT of 3.
- Sequence Number (16 bits) – used as indication for lost packets (or packets out of order).
  - Sequence Number is not used for playing the audio.

- Timestamp (32 bits) – Used for playing the audio.
- Some apps accept and play packets with timestamp of 0.
  - Some CISCO gateways do no permit it.
- SSRC – (32 bit) integer chosen randomly by each client once (when joining the session).
  - Each client should have a unique SSRC since it identifies it.
  - In a case when audio is sent from two different origins on the same RTP session, each should have its unique SSRC so the other side could identify it.
  - In case of a collision, RTCP BYE should be sent.

- CSRC list: 0 to 15 items, 32 bits each
- Not all members of the RTP header are used.
  - When can we take advantage of this fact ?
  - RTP or RTCP are not assigned any well known port number.
  - For RTP, port number must be even.
  - For RTCP, port number must be odd.
    - Usually, RTCP port = RTP port + 1
    - This was mandatory in the past, but it is not mandatory according to the last RTP spec.

SIP protocol

- SIP protocol (Session Initiation Protocol)
  - SIP is an application layer protocol
  - RFC 3261 (269 pages)

- SIP requests and SIP responses.
- Sip clients usually use UDP (port 5060).
  - However, also TCP can be used for SIP.
- Parameters for the SIP session (such as codecs and RTP port numbers) are negotiated by **SDP** (Session Description Protocol), which can be part of a SIP INVITE request.
  - Rfc4566
- When there is no common codec between two clients, we get **488** error - **Not acceptable here** response.

- The client sends a SIP Register request to a **registrar.**
- Example: a Register request was sent from **ekiga** Sip client to a Sip Registrar (**Kamailio**).
- The Sip Registrar returned **OK** (200).

- The **Use-Agent** header in the request is “ekiga”. (the client).
- The **Contact** includes the IP address of the SIP client.
- The **Expires** header tells for how long the registration is valid (3600 seconds, which is 1 hour).
  - In order to stay registered, a client must send another register request within an hour.

- INVITE – to ask the other side to start a call.
  - Usually there is an SDP as part of the INVITE request.
- BYE - end a call.
- REINVITE – for example, when you want to change the codec during a conversation.

- VoIP Clients for Linux:
- Ekiga (formerly GnomeMeething).
  - an open source project. (C++).
  - http://www.ekiga.net/
  - Maintained by Damien Sandras.
    - GNU General Public License (GPL)
- Provides a free SIP account.
- Support for these Audio Codecs:
  - Speex (8 KHz,16KHz), PCMU, PCMA, gsm, G726(16,24,32,40), ms-gsm, G722(16 Khz), ILBC.
- Support for these Video Codecs: H261.

- Eybeam (version 2.0) (there are free versions but it is NOT an open source project)
- http://www.counterpath.net/home.html
-

- A free version called **X-lite** (does not support advanced features, like g.729 codec).
  - Supports linux.
- A non-free version **(eyeBeam) for linux.**
  - Includes g.729 codec support.
    - Currently is not sold by countepath
    - In the past it was sold
      - but should be available again through 2010.

- **KPhone**
  - **http://sourceforge.net/projects/kphone/**
- **Skype** is offered for Linux since June 2004.
  - Freeware, but closed source.
  - Recently there were rumors that the company is working on an open source client for Linux.
    - Some say that this will be only for the UI layer.
    - There is no official info from skype about which protocols are used by skype.

Open source Servers

- Kamailio
  - http://www.kamailio.org/
  - There was an open source project called OpenSER
    - Was split into:
      - Kamailio
      - OpeSIPS (http://www.opensips.org/)
- Kamailio stores info in mysql DB.
- Adding users is simple.
  - kamctl add 1000 12345 (adding user 1000 with passwd 12345)
  - *kamctl db show subscriber* – show users.

Cellular phones

- **Currently, there are four big players in the market:**
  - **IPhone**
    - Officially entered the Israeli market in December 2009.
      - A while after it entered some other countries.
      - Based on MacOS.
      - Leopard is the final version of Mac OS X to support the PowerPC architecture as Snow Leopard (10.6) solely functions on Intel based Macs.
      - The OS on IPhone isn't a multitasking OS (as opposed to other cellulars).

- VoIP in IPhone;
- With the releae of iPad, in 27.1.10, apple also released updated iPhone SDK which **allows VoIP calls over 3G.**
- Taking in account that some cellular companies provide an unlimited 3G access, this should obviously be much cheaper to these users.

- **Nokia (Symbian)**
  - A known advantage of Symbian OS is long battery life.
  - VoIP: Recently, Skype released a beta version of Skype for some Nokia wireless phones.
    - Till then, there was Skype only for one or two Nokia phones.

- Seeking other markets:
  - August 2009: Nokia launches first Netbook, Nokia Booklet 3G.
    - Windows 7.
    - Intel Atom Chipset at 1.6 Ghz.
    - Note that there are now also ARM-based netbooks, for example, from Lenovo (which is good for linux fans: why?)
- February 2010:Symbian is released as Open Source
  - Prior to the intended date (June 2010).
  - First phones with the open source version of Symbian are due to beginning of 2011.
    - Eclipse Public License (EPL).
      - You don't have to share code.

- Are all Symbian packages released ?
- According to Symbian Foundation:
  - “All 108 packages containing the source code of the Symbian platform can now be downloaded from Symbian's developer web site".
  - Not according to Harald Welte's blog: (see in links below)
  - For example, no source code for these 2 packages:
    - **phonesrv** (implementing call engine, cell broadcast and SIM toolkit APIs)
    - **poc** (implementing push-to-talk).

Maemo

- **Maemo** is an operating system developed by Nokia for smartphones and Internet Tablets.
  - Maemo is based on Debian.
- You have a command line from an Xterm.
  - You can compile and build apps.
- Recently, Nokia released the N900 - a phone based on Linux (Maemo).
- Some saw in this step a sign that Nokia is moving to Linux.
- Is it indeed the case ?

- Keep in mind that Nokia is relatively not so popular in USA (Nokia gained much more popularity in Europe).

- N900 is the first tablet which is also a phone.
- There was a (second) Maemo developers summit of more than 400 participants, October 2009, Amsterdam.

- Maemo 5 uses a closed-source phone stack called the cellular service daemon (CSD).
- **Jun 2009 :**
  - Intel and Nokia jointly announce the oFono project.
    - Open source phone stack. (GPLv2)
    - http://ofono.org/
    - ofonod is a daemon controlled by D-Bus messages.
    - AT commands
      - over /dev/rfcomm0
      - Currently for GSM/UMTS
      - Future: CDMA/EVDO
    - Will probably not be included in Maemo 6.

Android

- Short history:
  - July 2005: Google acquired Android inc.
  - September 2008: Android 1.0 release.
    - Problems when developing VoIP apps when working with audio (soundcard).
  - April 2009: Android **1.5** (Cupcake).
  - September 2009: Android **1.6** (Donut).
  - October 2009: Android **2.0** (Eclair).

- Droid - Motorolla phone running Android 2.0.
  - The European version is called “Milestone Droid”.
  - Not shipped yet in Israel.
  - Has a physical keyboard.
  - A little heavier than Samsung galaxy (169 gram).
  - Milestone will be soon shipped in Israel (Calcalist, ynet). First in Cellcom, maybe by other service providers. It will probably be with Android 2.1
- Nexus One.
  - Android **2.1**
  - Stronger processor - **1 GHz** (Qualcomm QSD 8250)

- In Israel – Samsung Galaxy (cellcom)
  - Without physical keyboard.
  - Runs Android 1.5.
  - ARMv6-compatible processor
  - Besides phones, there are **netbooks** running Android (Acer and others) and **tablets** (ViewSonic and others).
  - Codecs: Packet Video.
- Also Sony Ericson entered the Android market with Xperia X10.
  - January 2010.

Android - contd

- Android git tree: http://android.git.kernel.org/
- Android does not have a glibc.
  - Instead, it has a special, open source libc ("Bionic").
  - Partial pthread support.
  - No SysV IPC support.
  - Instead, there is Android Binder IPC Driver
    - Based on OpenBinder (PalmOS)
    - Until 2.6.33 :
      - drivers/staging/android/binder.c
      - drivers/staging/android/binder.h
  - Smaller size than glibc (think of uLibc in embedded)

- Some criticized google for this:
- Matt Porter was heavily involved in the MIPS and PPC ports of Android.
- Matt gave a presentation titled “Android Mythbusters” in Embedded Linux Conference Europe , October 2009 (Grenoble, France)(see links).
- “The presentation shows how Google has simply thrown 5-10 years of Linux userspace evolution into the trashcan and re-implemented it partially for no reason.“
  - From Harald Welte's blog (see links).

- Also Matthew Garrett has a criticism about Android in his blog.
  - “Something like Binder **is pretty clearly not going upstream”.**
  - See links.

- Android was **removed** (December 2009) from the staging kernel tree (a git repository maintained by GregKH).
- GregKH is the maintainer of the staging tree.
  - GregKH came to a conclusion that these drivers can not get merged into the main kernel tree because they have dependencies on code that only lives in Google's kernel tree, causing it to fail to build in the kernel.org tree.

Staging: android: delete android drivers

So sad :(

...

15 files changed:

drivers/staging/android/Kconfig

drivers/staging/android/Makefile

drivers/staging/android/TODO

drivers/staging/android/binder.c

drivers/staging/android/binder.h

drivers/staging/android/logger.c

drivers/staging/android/logger.h

drivers/staging/android/lowmemorykiller.c

drivers/staging/android/ram_console.c

drivers/staging/android/timed_gpio.c

drivers/staging/android/timed_gpio.h

drivers/staging/android/timed_output.c

drivers/staging/android/timed_output.h

- DPI - Deep packet inspection.
  - Looks inside payload of packets for stings/regular expressions.
  - Works in Layer 2 through Layer 7.
  - China blocks VOIP traffic (in and out of China)
    - Skype seems to work without problem in China.
  - Vendors: Allot communications, others.
    - Vendors want to sell as much as possible; without VOIP, the market will be smaller.
    - There were some law suits in the USA
    - General opinion is that eventually the service provideres will not block VOIP. (Avi Patir, MIRS)

- **Moblin** - Intel.
- http://moblin.org/
  - Moblin stands for: 'mobile Linux',
  - 1 of April 2009.
    - Intel turned Moblin over to the Linux Foundation

Links

- **"Android Mythbusters"**
  - Matt Porter's presentation at the Embedded Linux Conference Europe , 2009.
- RTP book:

RTP: Audio and Video for the Internet

- By Colin Perkins
- Publisher: Addison Wesley
- Dual boot Maemo and Andorid:
  - http://www.youtube.com/watch?v=yri4qOfP8T0

- Intel® Integrated Performance Primitives (Intel® IPP)
- http://software.intel.com/en-us/intel-ipp/

- Removal of android from the staging tree:
  - http://www.kroah.com/log/linux/android-kernel-problems.html?seemore=y
- Voip in Iphone: the new sdk (27.1.10)
  http://www.downloadsquad.com/2010/01/27/apple-opens-up-voip-via-3g-on-iphone-and-maybe-even-ipad/

- Harald Welte's blog:
  - Symbian is Open Soruce - Really?
  - http://laforge.gnumonks.org/weblog/2010/02/05/#20100
- Matthew Garrett blog about Android
  - http://mjg59.livejournal.com/100221.html

Thank you!