

# OSADL License Compliance Audit

Till Jaeger - JBB Rechtsanwälte  
Armijn Hemel - Tjaldur Software Governance Solutions

December 4/5, 2014

# About OSADL

The Open Source Automation Development Lab (OSADL) is a registered cooperative under German law aimed at reducing cost for its members (43 regular members, 19 academic members, 5 sponsoring members) while solving shared problems:

- ▶ (real time) kernel development
- ▶ testing infrastructure
- ▶ research into safety critical systems certification
- ▶ legal issues

If you are interested in knowing more/joining we have brochures.

# Product fields

OSADL members are in fields that are different from for example consumer electronics (or making “Flappy Bird” clones):

- ▶ factory/building automation
- ▶ lasers
- ▶ robots
- ▶ washing machines (normal ones, but also for DNA)
- ▶ woodwork machines
- ▶ agriculture
- ▶ etc.

Many products are not flashy, but perform key roles in infrastructure.

# Considerations OSADL member companies have

Products have a long life. Some have a “return on investment” of 20 years. Support contracts often are (at least) as long. Facebook will be long gone by then!!

Software upgrades are frequently very hard, or almost impossible to do, or could be very risky, costly or even dangerous.

Many of these companies are therefore, by nature, risk averse.

Linux and Open Source helps them reduce this risk.

# OSADL license compliance audit

The license compliance audit (LCA) was developed because OSADL member companies requested it.

The LCA is limited to a single device/firmware revision combination.

Two components on the system are audited:

1. Linux kernel including external kernel modules
2. C library (glibc/uClibc), including how other components are linked to it

The audit is restricted to license compliance: detecting code theft/plagiarism, patent issues and license compatibility of user space programs are not part of the audit.

After successful completion of the audit the company gets a certificate.

# CERTIFICATE

Open Source License Compliance Audit



# Why did OSADL start this audit?

A few practical reasons:

- ▶ many OSADL companies are (originally) not software companies, a lot of knowledge is missing, yet open source is their (software) future
- ▶ wrong choices made now could have effects lasting decades

The actual goal for companies is not the certificate, but getting a “sanity check” and a better insight how well they are doing.

Also: no one had dared doing this before, so there was a challenge :-)

# Compliance audit setup

The audit is on site and lasts a full day. The company has one (or more) engineers involved in making the product available to answer any questions.

Furthermore we ask for:

- ▶ source code
- ▶ binary image
- ▶ documentation

The audit is done in a continuous dialogue and discussion between lawyer, audit engineer and company engineers to ensure knowledge transfer from the auditors to the company.



# Report

For each audit an extensive report with the conclusions of the audit is made. In the report problems found are described, plus steps that should be taken to remedy these problems.

# Questionnaire

Before we do an audit we ask companies to fill in a questionnaire:

- ▶ gives us some idea of what to expect
- ▶ ensures that we have all necessary information to actually do an audit
- ▶ forces companies to think about issues in advance

# Legal work done during the audit

1. license text review (compatibility, completeness)
2. contract & EULA review
3. written offer review
4. discussion about alternate ways to carry out license obligations

Many licenses require that license texts and author information are present. The review is to catch if licenses are missing or incomplete.

Contracts and EULA are checked for compatibility to see if they are compatible with GPL/LGPL.

The written offer (if any) is reviewed to see if it is correct.

# Technical work done during the audit

1. license scan
2. linking analysis
3. rebuild of kernel + C library/toolchain

## First: reducing the search space

Even though scope is limited it is still of work: recent Linux kernels have around 40,000 files.

Looking at each file would be very costly (at 1 second per file it would be 11+ hours), so we take a few shortcuts.

1. keep a database with all files in upstream `kernel.org` kernels
2. determine a checksum for each file in the kernel we audit
3. if the file can be found in the database we trust it and ignore it
4. look in more detail at files that can't be found in the database

For the Linux kernel this typically eliminates between 95% (but usually closer to 98%) of the files and allows us to focus on the *real* problems.

We think this is justified as the Linux kernel has already done due diligence on code it publishes. Similar: FSF (glibc) and SFC (uClibc).

# License scan

Each file that could not be found in the database (a few hundred files) are scanned with Nomos (part of FOSSology) and Ninka. This is to easily spot problematic files and prioritize work.

The license and copyright notices of each of these files is also verified manually.

# Linking analysis and kernel/C library rebuild

Using the Binary Analysis Tool (BAT) linking relationships between components are researched and visualised.

We also ask the company's engineers to rebuild the kernel and C library (or entire toolchain) on the spot and we compare the results with binaries on the device.

# Common issues found: technical

- ▶ missing source code
- ▶ incorrectly licensed source code



## Common issues found: legal

- ▶ contradicting statements in licenses and contracts
- ▶ LGPLv2.1 section 6 (requiring to allow reverse engineering and relinking)

## Common issues found: organisational

- ▶ tasks assigned to the wrong staff (who is responsible, who handles source code requests, etc.)
- ▶ purchase department does not take OSS licensing into account
- ▶ legal documents are not made or reviewed by the legal department, but by the development department (sometimes more than one)

# Current developments

We have started developing a *supplier* audit.

We hope to finish this in early 2015 and run first trials soon after.

## In conclusion: our experience

We have found this type of audit to be the most effective “tool” available to help achieve license compliance.

Although it is in essence a *product* audit we uncover quite a few *process* defects: many problems tend to be shared between multiple products.

If you want an audit, come talk to us! :-)

# Q&A