# AloisProposal

## ALOIS Proposal

ALOIS is a log collection and correlation software with reporting and alarming functionalities. It has been implemented by the Swiss company IMSEC for a customer about five years ago. GPL-licenced, implemented in Ruby and completely based on other OSS-licensed components, it was designed for the open source community right from the start. Now that the software has shown its functioning over several years in production with the one customer and one IMSEC-internal installation, it seems to be the right time to open it to a wider community.



## Abstract


ALOIS stands for "Advanced Logging and Intrusion Detection System" and is meant to be a fully implemented open source SIEM security information and event management) system.


## Proposal

While almost all other SIEM software, be it closed or open source, concentrate on the technological part of security monitoring, ALOIS is aimed to monitor the security of the content. It intends to be pro-active in the detection of potential loss, theft, mistaken modification or unauthorized access. ALOIS works on log messages and thus contains all the basic functionality of a conventional SIEM, as centralized collecting, normalizing, aggregation, analyzing and correlating of all log messages, as well as reporting all security related events. Therefore it can be used as any other SIEM.

ALOIS consists of five modules interacting to ensure a scaleable functionality of a SIEM:

- Insink is the message sink, which is the receiving entry point for all the different log messages into ALOIS. It is partly based on the syslog-ng software. Insink listens for messages (UDP), waits for messages (TCP), receives message collections (files, emails) and pre-filters them to prevent from message flow overload.
- Pumpy is the incoming FIFO buffer, implemented as a relational database tables. which contain the incoming original messages (in raw format). In a complex system setup, there may be several insink instances, e.g. for a group of hosts, for specific types of messages, or for high-avaliablity.
- Prisma contains logic to split up the text of log messages into separate fields, based on regular expressions. Actually, "prisma" is a set of "prismi", each one prisma for one type of log message (apache, cisco etc. Several prismi can be applied to the same message. This allows for stacked messages, i.e. forwarded log messages contained in compressed files contained in e-mail messages. The data retrieved form the log messages is stored in a database called Dobby. Due to prisma being written in Ruby, prismi can be applied interactively (when having system access).
- - Dobby is the central log database. It should be separated from the Pumpy database for availability and performance reasons. The current implementation is based on MySQL.
- - The Analyzer contains the two sub-systems Lizard and Reptor. Lizard is the analysis engine and user interface of ALOIS, implemented in Ruby on Rails using AJAX. It allows for interactive browsing through the collected data, exclusion/inclusion/selection of data, data

sorting, data filtering, creation of views, ad-hoc textual and graphical reporting. Reptor allows for automatic activation of views and comparison of these views' results to a predefined result (pattern matching). In case of mismatch, Reptor sends the result to predefined e-mail addresses.

Its modular design guarantees ALOIS to scale from little to large organizations. Since there exists a Debian package, it's easy to build a test system or even a productive system for small environments.

Although the software has been in productive use for a few years, there is still a lot of desired functionality missing. The plugability of new connected systems is given, but needs some revision. It is a given goal of the project to allow modules in other programming language. Furthermore, it has been discussed if parts of the existing implementation may be replaced with other proven open source software, e.g. the correlation engine or the web frontend. The other way round, it has been discussed that the filter creation engine would make a good tool for any kind of structured data, and thus could be separated from ALOIS and standardized as a stand-alone tool.

## Background

It's not simple to know what happens in a bigger network. There's a multitude of applications, services and appliances working together. Many of them provide some kind of events or state information. The network administrator needs to get hands on all of them. But they come in many different flavors and multiple canals. Therefore, it's hard to get the big picture. Furthermore, we have learned that it's impossible to protect a system against all malicious attacks and to keep all the possible faulty handling away. A monitoring of the systems to guarantee a pro-active handling is therefore needed..

Therefore, more and more organizations collect and analyze all logfiles in a centralized system, called a SIEM (security information and event management). The technology provides two major functions for security events from networks, systems and applications: log management and compliance reporting (SIM - security information management) and real-time monitoring and incident management (SEM - security event management).

## Rationale

Why another security information and event management system? It's true, there's already plenty of them. While the proprietary software is way too expensive for smaller to mid-sized companies, we find that the open source solutions are either too simple or not completely open. For example, behind each of the well known systems "OSSIM" and "Prelude", there is a company that either closes central functionality for its own business or has dual licensing and therefore asks the full copyright for all contributed code.

ALOIS is aimed to be totally free and open for all contributions. The openness provided for other programming languages is certainly proof of this. The plug-ability - yet to be further developed - is meant to guarantee that individual needs can be realized without stressing the whole system too much. In our opinion, the Linux kernel is a good example that this can work very well.

Since we are in accordance with "the Apache way", we would be very pleased if ALOIS could become part of the Apache community. In Addition, the Apache Logging Services would be a perfect home for the software. Furthermore, it's not the intention to compete with the already existing log viewer and analyzing tool "Chainsaw". Since Chainsaw is a relatively easy tool, it meets a rather different need. Nevertheless, if the two projects use synergies, both can profit.

## Initial Goals

When this project started ins 2005, there was no proven SIEM open source software and the commercial tools were way too expensive for the needed environment. Therefore, we decided together with a customer of ours to implement an open source SIEM tool from scratch. Now the software has run in a production environment for several years and has proven its functionality and reliabilty.

## Current Status

### Meritocracy

As already mentioned, ALOIS is already in production use in two organizations. All the code has been written by two persons of the same company in a paid employment relationship. It is obvious that this is way different from the open source approach within Apache. But nevertheless, the two developers have always worked as a team and the decisions were made in consensus whenever possible. But it is no secret, that these developers have to learn to behave in an open community. Understanding this potential problem, they already got support by a freelance consulter, who has the corresponding experience and knowledge.

### Community

Until today there is no real community, because the project hasn't been published officially, although it had been completely published on the web site for a couple of months (until a server relaunch). Convinced by the concept and design of the software, we are open and hope to reach many contributors and users. We think that it is realistic, because the SIEM issue has yet not been resolved in the OSS space.

### Core Developers

ALOIS was developed by Simon Hürliman and Flavio Pellanda, both employed by the company IMSEC. Concerning Design and Architecture, Marcus Holthaus, owner of IMSEC, gave his input as security specialist. Since the beginning of this year, Urs Lerch, a doctorate on the subject of commercial open source software development, supports the team with his knowledge. Simon Hülimann has left the company three years ago, but is still active in the OSS environment (although not for ALOIS). Current employee Daniel Lutz (a Debian Developer) has also contributed to the project.

### Alignment

Besides that we strongly believe in the "Apache way", we think that although that Apache hosts the Logging Services and different security projects, there is a gap when it comes to a superordinate security view. We therefore think it a good idea to add our SIEM project to the Apache repository. On the other side, Apache would become an even more complete software repository.

# Known Risks

## Orphaned products

Since the software is only maintained by employers of one company, there is a severe risk of being orphaned. But, on the one hand, the company has a sustained interest in keeping the project alive, because there are plans to offer services on top of ALOIS, and IMSEC uses the software for SIEM on their own systems. For this reason there exists a budget for the development and support of ALOIS. On the other hand, we believe that ALOIS is of great interest for other people and companies tied to IT security. Therefore, our step to the Apache incubator is also a step to a bigger community.

## Inexperience with Open Source

While ALOIS has always been licenced under the GPL, access to the source code, bug tracker and version control system has been restricted to internal users for most of the time. But the company has a strong believe in the open source movement and therefore engages its employees to take part in the community. Furthermore, it is also a strategic decision to build services on top of linux.

We understand that the Apache Incubator is a great opportunity for us to get assistance, when it comes to specific questions on the open source development. Even more, the company has created a part time position for the open source community work.

## Homogenous Developers

Although ALOIS has been developed by employees of only one company, there is a thorough openness. The company is designed to stay small and therefore works with several independent partners. Furthermore, its employees work in geographically different parts of the country. Therefore, it is no new experience for the developers to work in a distributed environment and argue rather than to command. Already today the employees are enforced to document all face-to-face communication in the internal wiki. Sketches are photographed and stored in the project's digital folder.

## Reliance on Salaried Developers

Until today all the development of ALOIS has been made in a paid emplyoment. Therefore we know that this brings a significant danger. Since it is our stated aim to encourage participation and recruit commiters, we hope to eliminate this risk as soon as possible. Furthermore, the employees of IMSEC are all open source enthusiasts and are in one way or another active in the community. Although we have no certainty, there is good indication that the current commiters would continue their work on ALOIS, even if they wouldn't be paid for it.

## Relationships with Other Apache Products

The Apache Logging Service would be a perfect home for ALOIS as a centralized logging collection and analyzing tool. Furthermore, we think that we could share part of the code with the Chainsaw subproject, since both need similar functionality in the web frontend. Since it is our statet aim to replace our own code with proofen open source libraries, we are open for any collaboration with other projects. For example, the replacement of the MySQL with a NoSQL database might be useful for performance reasons; therefore HBase is a good candidate.

## An Excessive Fascination with the Apache Brand

The Apache brand is in fact for its own a very good reason to join the Incubator. But much more our desire to become part of the Apache Incubator is our strong believe in open source software in general and in the "Apache way" in particular. We would love to learn from the experience and knowledge of the foundation's members and participants, which is an important part of the brand as well. The foundation has shown many times, that it has the processes and people to succeed in launching a project. We would be very proud to be part of this success story.

# Documentation

The documentation is rather weak and scattered. It has mainly been maintained on a wiki and is open to improvement. Since we are totally aware that this is a killer for a successfull open source project, we have already started an internal project with its own budget to improve this shortcomming. Once the project has been launched, writing a blog or open a forum are other possibilities we already thought of.

Furthermore, as the employees are used to work in a geographycally distributed environment, a lot of the internal communication happens in a chat. Thus, opening a new chat channel for the community is scheduled. (To document the discussions for all those who were off-line, we would send the logs daily to the mailing list.)

# Initial Source

Although the initial source comes from a project for a customer. it has an open source licence since the beginning. Therefore it doesn't have any propriatary code in it. A thorough revision before releasing it to a public repository is recommend and is also in planning.

The initial source will be a snapshot of the version control system, accompanied by a related debian package.

# Source and Intellectual Property Submission Plan

ALOIS is currently under a GPL licence. Since there are only two contributors so far, both from the same company, there is no problem to re-licence the code and contribute it to Apache. The commitment of the company's owner has been granted.

## External Dependencies

So far, no external dependencies are known. As mentioned before, a thorough revision of the codebase is in planning. There it can be controlled, that no other licence is affected by the code.

## Cryptography

ALOIS does not involve cryptographic code.

## Required Resources

### Mailing lists

The following mailing lists will be required:

- alois-private
- alois-dev
- alois-commits
- alois-users

### Subversion Directory

https://svn.apache.org/repos/asf/incubator/alois

### Issue Tracking

JIRA ALOIS (ALOIS)

### Other Resources

None.

## Initial Commiters

| NAME | EMAIL | AFFILIATION | CLA |
|---|---|---|---|
| Flavio Pellanda | flavio.pellanda at logintas dot ch | IMSEC | no |
| Urs Lerch | mail at ulerch dot net | IMSEC | yes |
| Daniel Lutz | daniel.lutz at logintas dot ch | IMSEC | no |
| Marcus Holthaus | marcus.holthaus at imsec dot ch | IMSEC | no |

## Sponsors

### Champion

Scott Deboy <sdeboy at apache dot org>

### Nominated Mentors

- Scott Deboy <sdeboy at apache dot org>
- Christian Grobmeier <grobmeier at apache dot org>

### Sponsoring Entity

The Incubator PMC (requested)