

Code Review of Voting Stimmunterlagen Offline

CLIENT **Federal Chancellery**
DATE **June 23, 2023**
VERSION **2.1**
GIT COMMIT **b0372fc**
STATUS **Final**

CLASSIFICATION **Public**
AUTHORS **Philippe Oechslin, Thomas Hofer**
DISTRIBUTION **Federal Chancellery**
MODIFICATIONS **Added analysis of trusted build**




Contents

1	Introduction	1
1.1	Context	1
1.2	Execution of the work	1
1.3	Executive summary	1
2	Analysis	2
3	Analysis of the code	4
3.1	Third party software	4
3.2	Analysis specific to the identified attacks	4
3.3	Generic analysis of the security of the code	5
3.4	Deployment	6
4	Recommendations	7
5	Conclusions	8



OS Objectif Sécurité SA
Route Cité-Ouest 19 - CH-1196 Gland
+41 22 364 85 70 - info@objectif-securite.ch



1 Introduction

1.1 Context

This report contains our review of a specific software used in the e-voting solution provided by Abraxax to the Canton of St-Gallen. The review was mandated by the Federal Chancellery as part of the examination process according to OEV Article 10 paragraph 1.

The reviewed software (Voting Stimmunterlagen Offline) is used after the content of the voting cards has been created. It transforms the raw data (names, codes, texts in XML format) to PDF files that can be sent to the printing office. It is a third-party software, which is not developed by Swiss Post.

1.2 Execution of the work

The initial review was carried out in the weeks 4 and 5 of 2023. We were given the full set of source code and the resources used to build the code, as well as the packaged version of the code that is delivered to the cantons.

An update of this report followed the publication of the source code in Spring 2023, with several points brought up in earlier versions having been resolved in the meantime.

1.3 Executive summary

The analysis of the results of the tests led us to the following conclusions:

No significant security issue: We found no evidence of code that would enable various attack scenarios that we imagined for the specific threat model of the environment in which the software is executed.

Overall design and code quality issues: While not an immediate threat to security, some code quality and architectural flaws complicate maintainability of the software and might lead to issues in the future.

Not all source code is published yet: Although most of the code is published, the source code of some specific libraries is not public yet.

2 Analysis

The role of the software is to convert the voting cards from xml format to PDF, for printing. The security of the e-voting systems depends on the fact that the content of the voting card is not modified or revealed.

The software is installed on a secured standalone laptop, operated by at least two persons. Data is exchanged by USB keys.

A manual review of a sample of voting cards is already in place, to detect visible defects in the cards.

The software uses the following assets:

- **Identity of voters:** the voter register,
- **Definition of the ballots:** questions, answers, candidates,
- **Codes:** initialisation key, return codes, confirmation and finalisation code,
- **Signature key:** used for signing the produced PDF documents,
- **Encryption key:** used for protecting the document during their transfer to the printing office.

Malicious operation by the software could result in the following classes of attack:

Attack 1: Adding or removing cards:

This is mitigated by the manual verification of the number of cards and reclamations of voters who do not receive a voting card.

Attack 2: Visible manipulation of voting cards to compromise individual verifiability: inverting the return codes for voting options, inverting the text of questions, exchanging the names of candidates.

This is mitigated by the manual review of a sample of cards.

Attack 3: Leaking information to the printing service by hiding it in the PDF files:

This does not have an impact, as the printing service has access to all the data (identity of voters, codes, ballot, encryption key), except for the signing key. Since the printing office is in charge of verifying the signature, being able to create fake signatures would not be an advantage.

Attack 4: Leaking information to an attacker by having it printed on a voting card:

The most efficient attack would be to leak the encryption key on a voting card, potentially hiding it by some steganographic method. The attacker could obtain a copy of the encrypted cards while they are transmitted from the canton to the printing office and then decrypt them when the voting card is printed and delivered by mail.

A more straight-forward attack would be leak codes by adding them to a card.

Adding the codes of a single other card would allow the attacker to break the secrecy of the vote cast with that card.

If the codes of a significant number of other cards can be added to one or few cards, this card could be used change the outcome of the vote (break the *vote correctness*).

The only effective mitigation against this class of attack is a review of the code. Publishing the code would allow for a much more thorough review of the code.

It is important to note that for these attacks to succeed, the attacker may have to compromise some elements of the untrusted part of the voting system. For example, to break vote secrecy, an attacker

who has obtained the codes of a victim and breaks into the voting server can compare the codes that are returned to the victim with the leaked codes. While compromising the voting server could be difficult, the trust model mandates that the system be safe, even if all untrusted parts have been compromised.

We have identified two types of attackers:

- **Internal attackers:** An attacker who injects malicious code into the software, before it is delivered to the canton.
- **External attackers:** An attacker who injects code into the software through data that is given to the code. The code could for example be added to the address field of a voter, before the voter registry is imported.

Any malicious action by the operators or the software can be ignored in the analysis of the software because the operators already have the capability to manipulate the voting cards without the help of the software. Additionally, the operators are subject to strong security rules (e.g. 4 eyes principle) as mandated in Number 3 of the OEV Annex.

3 Analysis of the code

3.1 Third party software

The software makes use of following third party software:

- SDelete (signed by Microsoft), securely deletes files on Windows
- Razor library, uses templates to transform JSON data to HTML

3.2 Analysis specific to the identified attacks

- ✓ We searched for the code lines that operate on the encryption and signature keys in all application components. The only components that use them are included in the dedicated `CryptoTool` package and executable. We did not find any uses of the keys beyond the expected decryption and signature features. (part of Attack 4).
- ✓ We analysed the way the code generates the voting cards. The input data is spread over several XML files, which are first aggregated into a single JSON model. The JSON model is then processed and broken down into individual voting cards in HTML format, which is then further transformed into PDF files. The HTML to PDF transformation is very straight-forward and does not perform any change on data. The input to JSON transformation relies on parsers for the dedicated files, and we did not find any evidence of improper mapping of data. The JSON to HTML transformation is driven by templates, relying on a third-party library (RazorLight). The templates are not part of the code, but we were provided with sample templates. Those templates we were given never mix data between voting cards (part of Attack 4).
- ✓ The templates we were given do not contain any logic to invert mappings in conditional scenarios, and therefore a randomized manual verification on known good datasets would be sufficient to detect errors, as long as the templates themselves have not been tampered with (Attacks 2 & 3).
- ✓ We analysed how the PDF documents are zipped and encrypted. The application has a preview feature for the generated PDFs. The files can then be downloaded, after having been encrypted. The code of the application uses the same data for preview and as input to the encryption tool. This seems to guarantee that the PDF files that are used for review are the same as the ones that are transmitted to the printing office (Attacks 1 & 2).
- ✓ We observed that the inputs of the software are XML files, which are handled by a standard XML parser, relying on well-defined XSD schemas. While the application relies on external libraries for which we did not get access to the source for the intermediary object representation of the data, the use of a standard parser and a well-defined XSD schema should be sufficient in thwarting most injection risks. We further confirmed that typical characters that could be used for injections (`!"#$%&'()*+,-./:;<=>?@[\\]^_`{|}~`) are properly processed. While this does not guarantee that effective code injection is impossible, it makes it very improbable.
- ✓ The process as a whole is streamlined by the client interface, which ensures continuity of the data handling. The front-end passes the expected elements to the expected command-line utilities.

3.3 Generic analysis of the security of the code

3.3.1 Security issues

Problem 1 – Outdated framework

The version of the .NET framework targeted for the build (`netcoreapp2.0`) is vulnerable and should be upgraded.

Recommendation

It is the auditors understanding that a migration is in progress. In the shorter term, setting the target version to `netcoreapp2.2` would already be an improvement since version 2.2, while out of support, has no known critical vulnerabilities.

Update - May 2023 The framework has been updated to version `net6.0`.

Problem 2 – Frontend dependencies vulnerable



Most previously vulnerable dependencies have been updated to versions that do not include known vulnerabilities. Only two dependencies with known vulnerabilities remain, and both of them are required for build only.

Dependency	Version
<code>minimatch</code>	pulled by <code>electron-builder @ 23.6</code>
<code>webpack</code>	pulled by <code>@angular-devkit/build-angular @ 14.2.10</code>

Recommendation

Dependencies should be kept up to date and regularly scanned for known vulnerabilities (e.g using `dependency-check`¹).

3.3.2 Code quality review

-  Coding style and good practices are generally consistent within the backend code.
-  Several coding style and software quality issues have been vastly improved since a previous audit.

¹ <https://owasp.org/www-project-dependency-check/>

3.4 Deployment

The code can be compiled in a reproducible way. Instructions are given for creating a docker in which the application can be built with a controlled set of dependencies. The SHA-256 hashes of all artifacts, before and after signature, are published with each release of the software.

Although the compilation is reproducible, the fact that the executables are signed – which is a good security practice – results in different hashes of executable files compiled at different times.

A trusted build ceremony is executed with the presence of a representative of the end-users of the software. The hashes documented during the ceremony can be used by the end-users to verify that the software they deploy is identical to the one that was compiled in the ceremony.

We noticed the following issues with the build ceremony:

Problem 3 – Build ceremony issues

- The executable code is signed using SHA-1 as digest function. Although there does not seem to exist a practical attack in this case, SHA-1 is deprecated as digest algorithm for signatures.
- The build process makes use of pre-compiled packages (NugGets). The build ceremony does not demonstrate that these packages were actually built from the published source code.
- The source code of some packages is not public yet.

Recommendation

A more modern hash function should be used for the signature, e.g. SHA-2 or SHA-3, all packages should be built from source code during the build ceremony and the code of all packages should be published.

4 Recommendations

Based on our analysis and tests, we can make the following recommendations:

Deployment:

- Compare the hashes of the software with the hashes obtained during the build ceremony.
- Have the changes in the code reviewed before each deployment. This includes the template files which may be updated for each election.

Operations:

- Review a sample of voting cards before sending them to the printing office. Verify that the texts and codes are identical to a known good source.

We understand that the recommendations regarding deployment and operations are already implemented.

Development:

- Publish the last parts of the source code that are not public yet.
- Adapt the build ceremony to include building of all packages.
- Require that the developers of the software use a secure coding standard that is similar to the one required from the developers of the e-voting software at Swiss Post. While many issues have been fixed since an earlier version of this report, coding style and code quality guidelines would still be an improvement to the project.

5 Conclusions

The source code that we reviewed seems to faithfully translate the received XML data into PDF files of voting cards.

We identified 4 types of attacks that could be mounted in the specific setup in which the application is used. Most of them can be easily excluded by reviewing the code. However, a manual review of a sample of the voting cards is recommended as a complement.

Almost all source code is published, and it is built in a trusted build ceremony. While there are still some issues to be fixed, the build ceremony will guarantee that the code is obtained from the original source code and from a controlled set of dependencies.

Any changes in the code should be reviewed and the dependencies must regularly be checked, to verify that they are authentic and up to date.

If these recommendations are applied, we see no explicit danger in using the software. We note however that the code could be simplified to make it easier to analyse and build independently. The general security of the code should be increased by applying coding standards similar to the ones used by the developers of the e-voting software.