



# Auditing Report

Hardening Blockchain Security with Formal Methods

FOR



zkemail.nr



Veridise Inc.  
February 25, 2025

► **Prepared For:**

Mach34  
<https://mach34.space/>

► **Prepared By:**

Tyler Diamond  
Ian Neal  
Benjamin Sepanski

► **Contact Us:**

[contact@veridise.com](mailto:contact@veridise.com)

► **Version History:**

|               |               |
|---------------|---------------|
| Feb. 25, 2025 | V2            |
| Dec. 03, 2024 | V1            |
| Nov. 26, 2024 | Initial Draft |

# Contents

|   |            |
|---|------------|
| <b>Contents</b>   | <b>iii</b> |
| <b>1 Executive Summary</b>  | <b>1</b>   |
| <b>2 Project Dashboard</b>  | <b>4</b>   |
| <b>3 Security Assessment Goals and Scope</b>                                  | <b>5</b>   |
| 3.1 Security Assessment Goals . . . . .                                       | 5          |
| 3.2 Security Assessment Methodology & Scope . . . . .                         | 5          |
| 3.3 Classification of Vulnerabilities . . . . .                               | 6          |
| <b>4 Vulnerability Report</b>   | <b>7</b>   |
| 4.1 Detailed Description of Issues . . . . .                                  | 8          |
| 4.1.1 V-ZIM-VUL-001: bh= may occur in other DKIM tags . . . . .               | 8          |
| 4.1.2 V-ZIM-VUL-002: redc field is manipulable . . . . .                      | 10         |
| 4.1.3 V-ZIM-VUL-003: Multiple nullifiers can be generated for a signature . . | 12         |
| 4.1.4 V-ZIM-VUL-004: Un-normalized signature/DKIM keys . . . . .              | 13         |
| 4.1.5 V-ZIM-VUL-005: Non-standard Email parsing . . . . .                     | 15         |
| 4.1.6 V-ZIM-VUL-006: First header value chars not validated . . . . .         | 17         |
| 4.1.7 V-ZIM-VUL-007: Duplicate code . . . . .                                 | 19         |
| 4.1.8 V-ZIM-VUL-008: Unused program constructs . . . . .                      | 21         |
| 4.1.9 V-ZIM-VUL-009: Missing/incorrect documentation . . . . .                | 22         |
| 4.1.10 V-ZIM-VUL-010: Elements may silently be skipped in hashing . . . . .   | 23         |
| 4.1.11 V-ZIM-VUL-011: Wrong value is compared to dkim header length . . . .   | 24         |
| 4.1.12 V-ZIM-VUL-012: Email nullifiers may leak information on the message .  | 25         |
| 4.1.13 V-ZIM-VUL-013: Ignored DKIM Tags . . . . .                             | 27         |
| 4.1.14 V-ZIM-VUL-014: Optimization opportunities . . . . .                    | 28         |
| <b>Glossary</b>   | <b>29</b>  |

From Nov. 11, 2024 to Nov. 25, 2024, Mach34 engaged Veridise to conduct a security assessment of their zkemail.nr. The security assessment covered the [zero-knowledge circuits](#) involved in validating the [DKIM](#) signature of an email. The circuit focuses on extracting the body hash of the email from the provided header and validating that a given [RSA](#) public key has signed the email. Veridise conducted the assessment over 6 person-weeks, with 3 security analysts reviewing the project over 2 weeks on commit 2f81196. Due to the [Noir](#) zero knowledge language exposing many utilities to developers, Veridise engineers also investigated some of the utilities' implementations invoked by Mach34's circuit.

**Project Summary.** The security assessment covered a [Noir](#) implementation of zkemail\*, as well as some associated TypeScript helper functions. The circuits use [DKIM](#) signatures as attestations to an email's validity, then extracts information about the email from the hashed headers (which also include a hash of the email body). This allows users to prove that they received an email from an entity associated with a certain DKIM key, then demonstrate application-specific properties hold over the email header or body without revealing sensitive information contained in the email (such as the sending/receiving parties or other sensitive information stored in the body of the email).

**Code Assessment.** The zkemail.nr developers provided the source code of the circuits for the code review. The source code appears to be mostly original code written by the zkemail.nr developers. It implements similar functionality to the original zkemail TypeScript library<sup>†</sup>. It contains some documentation in the form of READMEs and documentation comments on functions. To facilitate the Veridise security analysts understanding of the code, the zkemail.nr developers also wrote and shared documentation that provided more detail on the intentions of the project and design decisions they made.

The source code contained a test suite, which the Veridise security analysts noted provided both positive and negative tests for signature validation, and also tested the typescript library that generates inputs to the circuit.

**Summary of Issues Detected.** The security assessment uncovered 14 issues, 1 of which ([V-ZIM-VUL-001](#)) was assessed to be of high severity by the Veridise analysts. Specifically, a malicious user may be able to bypass the body hash parsing, allowing the user to insert a body hash under their control and thereby take control of the contents of the proven email message. The Veridise analysts also identified 1 medium-severity issue ([V-ZIM-VUL-002](#)), which allows attackers to manipulate the structure of the provided public key and potentially forge signatures, as well as 4 low-severity issues, 7 warnings, and 1 informational finding.

---

\* <https://github.com/zkemail>

† <https://github.com/zkemail/zk-email-verify/tree/d0e6f7ff9bc5723dbb8f4607b9f9892177b6cafb/packages/helpers>

All 14 issues have been acknowledged by the Mach34. Of the 14 acknowledged issues, Mach34 has fixed 9 issues. Mach34 does not plan to fix the other 5 acknowledged issues at this time, as much of the audited code has been changed by further development.

**Recommendations.** After conducting the assessment of the protocol, the security analysts had a few suggestions to improve zkemail.nr.

*Warnings and unconstrained functions.* Resolve compiler warnings so that newly introduced issues can be identified easier. Additionally, wrap all unconstrained calls in `unsafe` blocks. Finally, take care to verify if a function called from the standard library is unsafe before using it (see related issue [V-ZIM-VUL-002](#)).

*Avoiding type-casting.* Casting to `u32` (or other sized types) can be a dangerous operation, as it may silently truncate bits. When a value is expected to be of a certain size, consider constraining it to be of that size rather than truncating it. This will decrease the attack surface of the protocol.

*Full DKIM tag validation.* As described in [V-ZIM-VUL-013](#), several DKIM tags are ignored. This will allow validation of expired signatures, may cause issues for applications which receive simple-canonicalized headers, or could lead to compatibility problems if future versions of DKIM are released. While these are unlikely to be common errors, full spec compliance is important for trusted adoption of this library. All spec deviations should be carefully noted in the documentation. For example, a quoted-printable encoding is assumed, but not checked for.

*Redacted publication over private computation.* When writing applications using the zkemail library, the Veridise analysts recommend publishing redacted emails rather than proving statements about fully private emails. This allows for human intervention and review, without sacrificing the requirements privacy or authentication.

Using zkemail for financial applications requires a high degree of trust in the proper interpretation of an email. There are many potential threats to applications relying on this library. For example, DKIM keys are commonly associated with multiple entities from an organization. Automated accounts frequently use the same DKIM key as customer support accounts, meaning that great care must be taken to ensure the sender of an email is indeed the intended sender. Additionally, email is frequently unstructured, and rarely conforms to a specific structure that can be relied upon like an API. Emails sent from the same entity may vary based on a user account's settings/preferences conveyed to the entities (such as the account's time zone, language, or currency), or may change without notice at entity's discretion.

These various factors may lead to increased efforts in phishing to receive spurious emails signed by the correct DKIM keys, black-box testing configurations to find a vulnerable configuration, or other more advanced attacks such as compromising target email accounts. Keeping emails fully private, and only revealing application-specific information, will increase the attack surface and make attacks more difficult to identify.

*Hard-coded constants.* Consider defining constants for ASCII characters rather than using hard-coded hex-code. This will increase readability, and reduce the load of maintaining the repository.

**Disclaimer.** We hope that this report is informative but provide no warranty of any kind, explicit or implied. The contents of this report should not be construed as a complete guarantee that the system is secure in all dimensions. In no event shall Veridise or any of its employees be liable for any claim, damages or other liability, whether in an action of contract, tort or otherwise, arising from, out of or in connection with the results reported here.

**Table 2.1:** Application Summary.

| Name       | Version | Type | Platform |
|------------|---------|------|----------|
| zkemail.nr | 2f81196 | Noir | Aztec    |

**Table 2.2:** Engagement Summary.

| Dates                 | Method | Consultants Engaged | Level of Effort |
|-----------------------|--------|---------------------|-----------------|
| Nov. 11–Nov. 25, 2024 | Manual | 3                   | 6 person-weeks  |

**Table 2.3:** Vulnerability Summary.

| Name                          | Number | Acknowledged | Fixed |
|-------------------------------|--------|--------------|-------|
| Critical-Severity Issues      | 0      | 0            | 0     |
| High-Severity Issues          | 1      | 1            | 0     |
| Medium-Severity Issues        | 1      | 1            | 1     |
| Low-Severity Issues           | 4      | 4            | 2     |
| Warning-Severity Issues       | 7      | 7            | 5     |
| Informational-Severity Issues | 1      | 1            | 1     |
| TOTAL                         | 14     | 14           | 9     |

**Table 2.4:** Category Breakdown.

| Name                      | Number |
|---------------------------|--------|
| Data Validation           | 5      |
| Maintainability           | 4      |
| Logic Error               | 1      |
| Under-Constrained Circuit | 1      |
| Usability Issue           | 1      |
| Information Leakage       | 1      |
| Optimization              | 1      |



## 3.1 Security Assessment Goals

The engagement was scoped to provide a security assessment of zkemail.nr's circuits. During the assessment, the security analysts aimed to answer questions such as:

- ▶ Is the project compliant with common mail transfer protocols like RFCs 5321, 5322, and related superseded RFCs?
- ▶ Can valid emails be rejected by the protocol?
- ▶ Is the partial and interstitial SHA256 implemented correctly?
- ▶ Is the DKIM signature verification correctly implemented?
- ▶ Is each relevant tag from the DKIM signature properly parsed and validated?
- ▶ Are header fields properly validated and guaranteed to be what was signed?
- ▶ How are duplicate header fields handled?
- ▶ Are any common zero-knowledge vulnerabilities such as missing state checks, incorrect control-flow handling, or privacy leakage present?
- ▶ Is email address parsing correctly implemented? Can extra data be injected into the email?
- ▶ What types of attacks could usage of this system lead to (e.g. phishing, privacy leaks, etc.)?

## 3.2 Security Assessment Methodology & Scope

**Security Assessment Methodology.** To address the questions above, the security assessment involved a thorough review by human experts.

*Scope.* The scope of this security assessment is limited to the `lib/src` and `js/src` folders of the source code provided by the zkemail.nr developers, which contains the zero-knowledge circuit implementation of the zkemail.nr along with the typescript libraries to generate inputs for the circuits.

Specifically, the following files were in scope of the review:

1. `js/src/`
  - a) `index.ts`
  - b) `utils.ts`
2. `lib/src/`
  - a) `headers/body_hash.nr`
  - b) `headers/email_address.nr`
  - c) `headers/mod.nr`
  - d) `dkim.nr`
  - e) `lib.nr`
  - f) `macro.nr`



- g) `masking.nr`
- h) `partial_hash.nr`
- i) `remove_soft_line_breaks.nr`

*Methodology.* Veridise security analysts inspected the provided tests and read the `zkemail.nr` documentation. They then began a review of the code.

During the security assessment, the Veridise security analysts regularly met with the `zkemail.nr` developers to ask questions about the code.

### 3.3 Classification of Vulnerabilities

When Veridise security analysts discover a possible security vulnerability, they must estimate its severity by weighing its potential impact against the likelihood that a problem will arise.

The severity of a vulnerability is evaluated according to the Table 3.1.

**Table 3.1:** Severity Breakdown.

|             | Somewhat Bad | Bad     | Very Bad | Protocol Breaking |
|-------------|--------------|---------|----------|-------------------|
| Not Likely  | Info         | Warning | Low      | Medium            |
| Likely      | Warning      | Low     | Medium   | High              |
| Very Likely | Low          | Medium  | High     | Critical          |

The likelihood of a vulnerability is evaluated according to the Table 3.2.

**Table 3.2:** Likelihood Breakdown

|             |  |
|-------------|--|
| Not Likely  | A small set of users must make a specific mistake  |
| Likely      | Requires a complex series of steps by almost any user(s)<br>- OR -<br>Requires a small set of users to perform an action |
| Very Likely | Can be easily performed by almost anyone   |

The impact of a vulnerability is evaluated according to the Table 3.3:

**Table 3.3:** Impact Breakdown

|                   |   |
|-------------------|---|
| Somewhat Bad      | Inconveniences a small number of users and can be fixed by the user   |
| Bad               | Affects a large number of people and can be fixed by the user<br>- OR -<br>Affects a very small number of people and requires aid to fix                                      |
| Very Bad          | Affects a large number of people and requires aid to fix<br>- OR -<br>Disrupts the intended behavior of the protocol for a small group of users through no fault of their own |
| Protocol Breaking | Disrupts the intended behavior of the protocol for a large group of users through no fault of their own   |



This section presents the vulnerabilities found during the security assessment. For each issue found, the type of the issue, its severity, location in the code base, and its current status (i.e., acknowledged, fixed, etc.) is specified. Table 4.1 summarizes the issues discovered:

**Table 4.1:** Summary of Discovered Vulnerabilities.

| ID            | Description                                      | Severity | Status       |
|---------------|--|----------|--------------|
| V-ZIM-VUL-001 | bh= may occur in other DKIM tags                 | High     | Acknowledged |
| V-ZIM-VUL-002 | redc field is manipulable                        | Medium   | Fixed        |
| V-ZIM-VUL-003 | Multiple nullifiers can be generated for a . . . | Low      | Acknowledged |
| V-ZIM-VUL-004 | Un-normalized signature/DKIM keys                | Low      | Fixed        |
| V-ZIM-VUL-005 | Non-standard Email parsing                       | Low      | Acknowledged |
| V-ZIM-VUL-006 | First header value chars not validated           | Low      | Fixed        |
| V-ZIM-VUL-007 | Duplicate code                                   | Warning  | Fixed        |
| V-ZIM-VUL-008 | Unused program constructs                        | Warning  | Fixed        |
| V-ZIM-VUL-009 | Missing/incorrect documentation                  | Warning  | Fixed        |
| V-ZIM-VUL-010 | Elements may silently be skipped in hashing      | Warning  | Fixed        |
| V-ZIM-VUL-011 | Wrong value is compared to dkim header . . .     | Warning  | Fixed        |
| V-ZIM-VUL-012 | Email nullifiers may leak information on . . .   | Warning  | Acknowledged |
| V-ZIM-VUL-013 | Ignored DKIM Tags                                | Warning  | Acknowledged |
| V-ZIM-VUL-014 | Optimization opportunities                       | Info     | Fixed        |

## 4.1 Detailed Description of Issues

### 4.1.1 V-ZIM-VUL-001: bh= may occur in other DKIM tags

|                         |  |               |              |
|-------------------------|--|---------------|--------------|
| <b>Severity</b>         | High   | <b>Commit</b> | 2f81196      |
| <b>Type</b>             | Data Validation  | <b>Status</b> | Acknowledged |
| <b>File(s)</b>          | lib/src/headers/body_hash.nr   |               |              |
| <b>Location(s)</b>      | get_body_hash()  |               |              |
| <b>Confirmed Fix At</b> | <a href="https://github.com/zkemail/zkemail.nr/pull/23/">https://github.com/zkemail/zkemail.nr/pull/23/</a> ,<br><a href="https://github.com/zkemail/zkemail.nr/pull/38/">https://github.com/zkemail/zkemail.nr/pull/38/</a> |               |              |

The `get_body_hash()` function obtains a sequence of length `BODY_HASH_BASE64_LENGTH` prefixed by a `bh=`, constrained to be within the `dkim-signature` header field. However, no check is performed to ensure that this sequence does not occur as part of another tag-value within the `dkim-signature` header-field.

Other tags may be user-controlled, leading to extraction of a user-controlled body-hash.

**Impact** If users create strange-looking email header field values (e.g. `bh=7xQMDuoVVU4m0W0WRVsrVXMeGSIASsnucK9dJsrc+vU=@domain.com`) or include the body-hash in the subject line when the mail-server is including the subject in the `z=` tag, the user may successfully be able to receive a `dkim-signature` header-value which contains *two* `bh=...` sequences: the actual body-hash and a fake body-hash.

Once a user has done this, they have full control over the contents in the proven message.

**Recommendation** Check that the extracted `bh=` tag is either the first tag, or preceded by a semicolon. See Section 3.2 in RFC 6376.

**Developer Response** We changed the code to

1. add a check for prefix of `;" bh=` instead of `"bh=` when accessing body hash from `dkim-header` field.
2. add a check for suffix `;"` (no suffix check previously) when accessing body hash from `dkim-header` field.

We also included a test with crafted `dkim-signature` field that should otherwise have passed old `get_body_hash` check but now fails.

**Updated Veridise Response** This will remove the issue. However, in the case that the very *first* DKIM-signature is `bh=`, it will fail. This may deny service to valid e-mails.

The case of `bh=` being the first DKIM tag should be handled.

**Updated Developer Response** We now also check for cases where `bh` is preceded by a colon.

**Updated Veridise Response** Colons can appear in tag values, so the `body_hash_index` must be validated to be at the start of the DKIM-signature when accepting : `bh=` as a prefix.

### 4.1.2 V-ZIM-VUL-002: redc field is manipulable

|                         |   |               |         |
|-------------------------|---|---------------|---------|
| <b>Severity</b>         | Medium  | <b>Commit</b> | 2f81196 |
| <b>Type</b>             | Data Validation   | <b>Status</b> | Fixed   |
| <b>File(s)</b>          | lib/src/dkim.nr   |               |         |
| <b>Location(s)</b>      | verify_dkim_signature()   |               |         |
| <b>Confirmed Fix At</b> | <a href="https://github.com/zkemail/zkemail.nr/pull/37">https://github.com/zkemail/zkemail.nr/pull/37</a> |               |         |

The `verify_dkim_signature()` function is a method attached to an `RSAPubkey` parameter representing the RSA public key to verify the DKIM signature against. This struct has 2 fields: the `modulus` and the `redc`. The latter parameter should have a specific relation to the `modulus` as defined [here](#). As an example, Z-imburse's `verify_linode_billing_receipt()` passes in the `RSAPubkey` as is:

```

1 pub fn verify_linode_billing_receipt(
2     header: BoundedVec<u8, MAX_LINODE_EMAIL_HEADER_LENGTH>,
3     pubkey: RSAPubkey<KEY_LIMBS_2048>,
4     signature: [Field; KEY_LIMBS_2048],
5     ...
6 ) -> [Field; 3] {
7     ...
8     pubkey.verify_dkim_signature(header, signature);
9     ...
10 }

```

#### Snippet 4.1: Snippet from `verify_linode_billing_receipt()`

The only verification of the `pubkey`'s legitimacy is in the `z_imburse_registry`, which confirms the `modulus` hash is registered. There is no verification that the `redc` parameter has the proper relation to `modulus`. Therefore, an attacker may be able to manipulate the parameter in order to find data that a forged signature can be created for.

**Impact** An attacker may be able to forge signatures against data that was not legitimately signed. Due to time constraints, Veridise analysts did not fully exploit this attack. However, it very clearly causes the implementation to diverge from RSA, and likely may be converted into full signature forgery given an additional few days.

**Recommendation** Verify that the `redc = floor((1 << (2 * Params::modulus_bits())) / modulus)`.

**Developer Response** We compromise between efficiency and security by including the `redc` limbs in the `dkim` key hash. Now the verifier must (once) compute additional constraints around `redc = floor((1 << (2 * Params::modulus_bits())) / modulus)`, however the prover would not need to do this every time. This also added a couple thousand constraints to the hash, which we mitigated by composing two 120-bit limbs into one 240-bit number.

**Updated Veridise Response** The provided `hash()` function is not injective. A 120-bit decomposition is only injective if both elements are known to be 120 bits. Instead, two fields are used.

This could potentially allow someone to use a different modulus.

**Updated Developer Response** We added a function which validates the `redc` and `modulus` are both reduced in `zkemail`.

### 4.1.3 V-ZIM-VUL-003: Multiple nullifiers can be generated for a signature

|                         |             |               |                    |
|-------------------------|-------------|---------------|--------------------|
| <b>Severity</b>         | Low         | <b>Commit</b> | 2f81196            |
| <b>Type</b>             | Logic Error | <b>Status</b> | Acknowledged       |
| <b>File(s)</b>          |             |               | lib/src/lib.nr     |
| <b>Location(s)</b>      |             |               | standard_outputs() |
| <b>Confirmed Fix At</b> |             |               | N/A                |

The `standard_outputs()` function calculates an `email_nullifier` as one of its outputs. It does so by simply taking the pedersen hash of the passed signature, as seen below:

```
1 let email_nullifier = pedersen_hash(signature);
```

#### Snippet 4.2: Snippet from lib/src/lib.nr:standard\_outputs()

The signature is composed of `KEY_BYTE_LENGTH` Fields which should be represented by 120-bit limbs. However, given that the `Field` type is much larger, one can simply add the RSA public key modulus to the signature to generate a signature that is distinct in its hash, but represents the same signed message. Therefore, one can generate a large amount of signatures from a single legitimately signed message. A proof-of-concept is provided below

```
1 #[test]
2 fn test_dkim_signature_add() {
3     let mut sig = EmailLarge::SIGNATURE;
4     let pubkey = EmailLarge::PUBKEY;
5     for i in 0..KEY_LIMBS_2048 {
6         sig[i] += pubkey.modulus[i];
7     }
8     pubkey.verify_dkim_signature(EmailLarge::HEADER, sig);
9 }
```

#### Snippet 4.3: A proof-of-concept of duplicating a signature, ran in lib/src/tests/mod.nr

**Impact** If the `email_nullifier` is used as a nullifier in consumers of this library, then one can generate a plethora of nullifiers for the same email.

**Recommendation** Reduce the signature by the public key's modulus before hashing it.

**Developer Response** The developers acknowledge the issue but have not provided a fix as certain signatures cause underflows in the `BigNumber` library used by the RSA verification library.

#### 4.1.4 V-ZIM-VUL-004: Un-normalized signature/DKIM keys

|                  |   |        |         |
|------------------|---|--------|---------|
| Severity         | Low   | Commit | 2f81196 |
| Type             | Under-Constrained C...  | Status | Fixed   |
| File(s)          | lib/src/dkim.nr   |        |         |
| Location(s)      | verify_dkim_signature()   |        |         |
| Confirmed Fix At | <a href="https://github.com/zkemail/zkemail.nr/pull/26">https://github.com/zkemail/zkemail.nr/pull/26</a> |        |         |

As mentioned in V-ZIM-VUL-003, signatures are represented as an array of limbs. The `verify_dkim_signature()` (shown below) operates on the signature represented as a `BigNum`.

```

1 // instantiate BN for the signature
2 let signature: BN1024 = BigNum::from_array(signature);
3
4 // verify the DKIM signature over the header
5 let rsa: RSA1024 = RSA {};
6 assert(rsa.verify_sha256_pkcs1v15(pubkey, header_hash, signature, RSA_EXPONENT));

```

**Snippet 4.4:** Snippet from `verify_dkim_signature()`.

The internal implementation of `BigNum` assumes the limbs have been range-checked to 120-bits so that no intermediate field multiplications inside the `BigNum.mul` methods used in `verify_sha256_pkcs1v15` overflow. However, this is never validated by the developers.

This means that attackers may produce additional signatures by using un-normalized representations of a DKIM signature. The below proof of concept can be generalized to a large class of alternate signatures

```

1 #[test]
2 fn test_dkim_signature_unnormalized() {
3     let mut sig = EmailLarge::SIGNATURE;
4     let pubkey = EmailLarge::PUBKEY;
5     let delta = 1;
6     sig[0] += delta * 0x100000000000000000000000000000000 ;
7     sig[1] -= delta;
8     pubkey.verify_dkim_signature(EmailLarge::HEADER, sig);
9 }

```

**Snippet 4.5:** Adding and subtracting between limbs still produces a valid signature. This test can be run from `lib/src/tests/mod.nr`.

**Impact** From a single signature, attackers may produce multiple fresh signatures which pass verification. Since pedersen hashes are used to commit to signatures in the repository, each of these malformed signatures will produce a different commitment, appearing to be a new valid signature in cases where the signature is kept as a private witness.

Additionally, attackers may attempt to leverage overflows in the native Noir field during RSA verification. While the Veridise analysts were unable to exploit this in the time allotted, a dedicated attacker may be able to leverage the field overflows to produce valid signatures without knowledge of the private key.

**Recommendation** Constrain the signature to have normalized limbs.



**Developer Response** The same signature can be represented multiple different ways without range checking the limbs of the signature. This fix simply uses `signature.validate_in_range()` to ensure each limb is checked. It additionally uses the prescribed unit test with an expectation of failure to demonstrate remediation of this issue.

#### 4.1.5 V-ZIM-VUL-005: Non-standard Email parsing

|                         |                 |               |                       |
|-------------------------|-----------------|---------------|-----------------------|
| <b>Severity</b>         | Low             | <b>Commit</b> | 2f81196               |
| <b>Type</b>             | Data Validation | <b>Status</b> | Acknowledged          |
| <b>File(s)</b>          |                 |               | See issue description |
| <b>Location(s)</b>      |                 |               | See issue description |
| <b>Confirmed Fix At</b> |                 |               | N/A                   |

Email addresses are frequently thought of as a sequence of alpha-numeric characters (with a few allowed exceptions like -, ., and \_) followed by an @, and then a domain. However, there are well-established standards for accepted forms of email addresses, including

1. The Address Specification in [Section 3.4 of IETF RFC 5322](#) (superseding RFCs 2822 and 822).
2. [IETF RFC 6531](#), which defines an extension for internationalized email.

Common implementations of email address parsers (e.g. implementations in [Rust](#) or [python](#)) expressly conform to one or more of the above RFCs. However, zk-email diverges in some cases from these specifications.

1. `js/src/utlils.ts`:
  - a) `getAddressHeaderSequence()`: This function restricts emails to include any sequence of "valid" characters (`[a-zA-Z0-9._%+-]`) followed by an @, and a similarly defined domain specification. This will wrongly accept some names (e.g. those with a repeated "." character) and wrongly deny others (e.g. those including a #).
  - b) `getHeaderSequence()`: Some headers may appear more than once. See the table just above [Section 3.6.1 in RFC 5322](#).
2. `lib/src/lib.nr`:
  - a) Some commonly used email characters, such as an under-score, are disallowed.
3. `lib/src/headers/email_address.nr`:
  - a) Strings without an @ may be accepted as valid email addresses. In general, there is no guarantee that the email addresses are valid.

**Impact** Some library users may be unable to prove desired properties about valid emails.

Other use cases may be able to choose invalid emails which pass verification but break important invariants of an application.

**Recommendation** In all JavaScript/TypeScript code, consider using an established and well-used email library.

Explicitly choose a specification to conform to, or clearly indicate the deviations from the specification in the documentation. This may be especially important for users who need to ensure their email addresses are valid.

**Developer Response** We acknowledge that this is an issue, however email address parsing currently is incredibly primitive and we do not intend on employing it in the future as we cannot parse cc's or multiple recipients easily in the current paradigm. Instead, we will integrate ZK-Regex which is quite capable of making complex queries over email addresses.

We acknowledge these issues and intend to solve them in the near future, but are neglecting to remedy them immediately as this technique will be deprecated.

#### 4.1.6 V-ZIM-VUL-006: First header value chars not validated

|                         |   |               |         |
|-------------------------|---|---------------|---------|
| <b>Severity</b>         | Low   | <b>Commit</b> | 2f81196 |
| <b>Type</b>             | Data Validation   | <b>Status</b> | Fixed   |
| <b>File(s)</b>          | lib/src/headers/mod.nr  |               |         |
| <b>Location(s)</b>      | constrain_header_field()  |               |         |
| <b>Confirmed Fix At</b> | <a href="https://github.com/zkemail/zkemail.nr/pull/25/">https://github.com/zkemail/zkemail.nr/pull/25/</a> |               |         |

The `constrain_header_field()` function checks that no CRLF sequences occur within the header value to prevent a malicious prover from including subsequent header field name/value pairs as part of the header value.

Unfortunately, as shown in the below snippet, this check is performed starting from `start_index+i`, where `i` starts from `HEADER_FIELD_NAME_LENGTH+1`. This means that the check is not applied to the field value characters at indices `HEADER_FIELD_NAME_LENGTH + 1, ..., 2 * HEADER_FIELD_NAME_LENGTH + 1`. As a consequence, a malicious prover may be able to include extra data in a header-value if the header field name is longer than the header field value.

```

1 let start_index = header_field_sequence.index + HEADER_FIELD_NAME_LENGTH + 1;
2 for i in (HEADER_FIELD_NAME_LENGTH + 1)..MAX_HEADER_FIELD_LENGTH {
3   // is it safe enough to cut this constraint cost in half by not checking lf? i
   think so
4   let index = start_index + i;
5   if (index < header_field_sequence.index + header_field_sequence.length) {
6     assert(
7       header.get_unchecked(index) != CR,
8       "Header field must not contain newlines"
9     );
10  }
11 }

```

**Snippet 4.6:** Snippet from `constrain_header_field()`

**Impact** This can be very important to ensure that user-controlled inputs (such as the subject line, which may be manipulable through phishing) are not included in a header for later exploitation.

For example, if an application is extracting subject from the below email,

```

1 subject:t
2 from:longstring

```

an attacker could bypass the newlines check. Since `HEADER_FIELD_NAME_LENGTH=7` for subject, the below malicious input

```

1 index = 0
2 length = 10

```

leads the application to compute

```
1 start_index = 0 + 7 + 1 = 8
2 index = start_index + i = 8 + 8 = 16
```

which is the index of the colon in `from:longstring`. If the user can control `longstring`, they may then lie to the application by proving facts about `longstring` instead of the actual subject (`t`).

Luckily, this only works if the header value is shorter than the header field-name, which is not commonly the case for any of `from`, `to`, or `dkim-signature`.

**Recommendation** Loop over the range starting from 0.

**Developer Response** The developers acknowledge the issue and now loop starting at index 0.

#### 4.1.7 V-ZIM-VUL-007: Duplicate code

|                  |   |        |         |
|------------------|---|--------|---------|
| Severity         | Warning   | Commit | 2f81196 |
| Type             | Maintainability   | Status | Fixed   |
| File(s)          | See issue description   |        |         |
| Location(s)      | See issue description   |        |         |
| Confirmed Fix At | <a href="https://github.com/zkemail/zkemail.nr/pull/29">https://github.com/zkemail/zkemail.nr/pull/29</a> |        |         |

Duplicate code or constraints were identified, and should be de-duplicated:

1. `lib/src/headers/mod.nr`:
  - a) The `end_index` variable is defined as the addition of `index` and `length` of the sequence instead of using the `end_index()` method.
    - i. Note that there are a few locations where this is done.
  - b) The `header_field_sequence` `index+length` is checked to be less than `header.len()`, and this check is repeated after `end_index` is defined.
  - c) The `constrain_header_field_detect_last_angle_bracket()` and `constrain_header_field()` functions share most of their implementation. Consider refactoring out the common logic into an internal function used by both implementations.
2. `lib/src/remove_soft_line_breaks.nr`:
  - a) In `find_zeros()` the `crate::CR` and `crate::LF` constants should be used instead of hard-coding the `0x0D` and `0x0A` values.
3. `lib/src/lib.nr`:
  - a) The `verify_email()` function asserts `header.len() <= MAX_EMAIL_HEADER_LENGTH`. This is already implicitly constrained due to `header` being a `BoundedVec`.
4. `lib/src/partial_hash.nr`:
  - a) The `partial_sha256_var_start`, `partial_sha256_var_interstitial`, and `partial_sha256_var_end` functions all duplicate a large amount of code. However, since this code is primarily copied from the Aztec standard library and de-duplicating this code would require significant refactoring, we do not find that de-duplicating code in these functions is strictly necessary.

**Impact** Duplicate code can make maintaining the codebase more difficult, and bugs may be introduced in the future if the implementations were to diverge. Additionally, duplicate constraints unnecessarily increase the circuit size.

**Recommendation** Refactor to remove the duplicated code.

**Developer Response** We made the suggested changes minus partial hash due to it requiring extensive changes:

1. `lib/src/headers/mod.nr`:
  - a) Remove unnecessary calculations for header `end_index`.
  - b) Remove duplicated checks on header `end_index`.
  - c) Abstract away duplicate logic for `constrain_header_field_detect_last_angle_bracket()` and `constrain_header_field()`.
2. `lib/src/remove_soft_line_breaks.nr`:
  - a) Use global variables for CR and LF in `find_zeroes()` comparison.
3. `lib/src/lib.nr`:
  - a) Remove redundant header length check since header bounds are constrained via `BoundedVec`.
4. `lib/src/partial_hash.nr`:
  - a) Leave code deduplication for partial hashing logic due it requiring large changes.

#### 4.1.8 V-ZIM-VUL-008: Unused program constructs

|                         |   |               |         |
|-------------------------|---|---------------|---------|
| <b>Severity</b>         | Warning   | <b>Commit</b> | 2f81196 |
| <b>Type</b>             | Maintainability   | <b>Status</b> | Fixed   |
| <b>File(s)</b>          | See issue description   |               |         |
| <b>Location(s)</b>      | See issue description   |               |         |
| <b>Confirmed Fix At</b> | <a href="https://github.com/zkemail/zkemail.nr/pull/32">https://github.com/zkemail/zkemail.nr/pull/32</a> |               |         |

**Description** The following program constructs are unused:

1. `lib/src/lib.nr`:
  - a) The `standard_outputs()` function does not appear to be used outside of the examples.
2. `lib/src/partial_hash.nr`:
  - a) In `msg_u8_to_u32()` the `msg32` array is hard-coded to be of length 16. This depends on `BLOCK_SIZE` always being 64. Instead of leaving the `BLOCK_SIZE` unused, meta-programming should be utilized to generate the length as `BLOCK_SIZE/4`, and assert the `BLOCK_SIZE` is a multiple of 4. Similar comments apply to the following functions:
    - i. `partial_sha256_var_end()`.
    - ii. `verify_msg_block()`.
  - b) The global `ZERO` defined in this file is unused.
  - c) The global `DATA` should be defined within `mod tests`.

**Impact** These constructs may become out of sync with the rest of the project, leading to errors if used in the future.

#### Developer Response

1. `lib/src/lib.nr`:
  - a) `standard_outputs()` was removed in prior PR.
2. `lib/src/partial_hash.nr`:
  - a) `msg_u8_to_u32()` was left unchanged as we feel allowing a variable block size is unnecessary given that blocks will always be 64 bytes in size for sha256.
  - b) global `ZERO` was taken out in prior PR.
  - c) Global `DATA` moved into `mod tests`.



#### 4.1.9 V-ZIM-VUL-009: Missing/incorrect documentation

|                         |   |                       |         |
|-------------------------|---|-----------------------|---------|
| <b>Severity</b>         | Warning   | <b>Commit</b>         | 2f81196 |
| <b>Type</b>             | Maintainability   | <b>Status</b>         | Fixed   |
| <b>File(s)</b>          |   | See issue description |         |
| <b>Location(s)</b>      |   | See issue description |         |
| <b>Confirmed Fix At</b> | <a href="https://github.com/zkemail/zkemail.nr/pull/30">https://github.com/zkemail/zkemail.nr/pull/30</a> |                       |         |

The following items have incorrect documentation or are lacking documentation:

1. `js/src/util.ts`
  - a) `makeEmailAddressCharTable()`:
    - i. This function contains an extraneous `console.log()`.
    - ii. Consider adding documentation noting that this function is used to generate Noir code, and link to the code location.
    - iii. The variable `proceedingChars` contains a minor typo. It should be defined as `proceedingChars`.
  - b) `u8tToU32()`: Consider including the endianness of the output U32s in the documentation.
  - c) The end of the file contains commented out dead code.
2. `lib/src/headers/body_hash.nr`:
  - a) `get_body_hash_unsafe()`: Consider linking to the implementation of `base64_encode` as explanation for the counterintuitive function naming. For example, [this link](#) may be a useful reference for readers.
3. `lib/src/headers/mod.nr`:
  - a) `constrain_header_field_detect_last_angle_bracket()`: The doc comment of this function has a typo, writing `conrain` instead of `constrain`.
4. `lib/src/lib.nr`:
  - a) `MAX_EMAIL_ADDRESS_LENGTH`: Consider linking to [RFC 5321](#) as an explanation for how this maximum length was derived.
5. `lib/src/partial_hash.nr`:
  - a) `partial_sha256_var_start()` details that `N` is the maximum length of the message to hash instead of the actual length. Additionally, `message_size` does not exist for this function.

**Impact** Missing or incorrect documentation may lead to misuse of the library down the road.

**Recommendation** Fix the documentation for the above functions.

**Developer Response** We made the recommended changes.

#### 4.1.10 V-ZIM-VUL-010: Elements may silently be skipped in hashing

|                         |   |               |         |
|-------------------------|---|---------------|---------|
| <b>Severity</b>         | Warning   | <b>Commit</b> | 2f81196 |
| <b>Type</b>             | Usability Issue   | <b>Status</b> | Fixed   |
| <b>File(s)</b>          | lib/src/partial_hash.nr   |               |         |
| <b>Location(s)</b>      | partial_sha256_var_start()  |               |         |
| <b>Confirmed Fix At</b> | <a href="https://github.com/zkemail/zkemail.nr/pull/28">https://github.com/zkemail/zkemail.nr/pull/28</a> |               |         |

In `partial_sha256_var_start`, the `num_blocks` to hash is defined as  $N / \text{BLOCK\_SIZE}$ . When  $N \% \text{BLOCK\_SIZE} \neq 0$ , the leftover bytes at the end of the `msg` are not mixed into the internal hashing state, and therefore are simply ignored. There is no indication that this is occurring.

**Impact** The outputted intermediate hash state will not contain the leftover bytes, and the caller of this function may expect that these bytes were included in the calculation. Given the intended usage of this function, any such errors are likely to be found by testing, but it should still be fixed.

**Recommendation** Fail if  $N \% \text{BLOCK\_SIZE} \neq 0$ , or make it explicitly clear that the caller must make this check.

**Developer Response** Instead of reverting in the case that  $N \% \text{BLOCK\_SIZE} \neq 0$ , we will allow users the flexibility to work around this and provide a warning message above the function notifying them of this issue

#### 4.1.11 V-ZIM-VUL-011: Wrong value is compared to dkim header length

|                         |   |               |         |
|-------------------------|---|---------------|---------|
| <b>Severity</b>         | Warning   | <b>Commit</b> | 2f81196 |
| <b>Type</b>             | Maintainability   | <b>Status</b> | Fixed   |
| <b>File(s)</b>          | src/headers/body_hash.nr  |               |         |
| <b>Location(s)</b>      | get_body_hash()   |               |         |
| <b>Confirmed Fix At</b> | <a href="https://github.com/zkemail/zkemail.nr/pull/27">https://github.com/zkemail/zkemail.nr/pull/27</a> |               |         |

In `get_body_hash()` the `body_hash_index` represents the location of where the body hash field body starts in the header. The header will be read starting at this value for `BODY_HASH_BASE64_LENGTH` bytes. In order to ensure that the bytes read are within the bound of the body hash header field, the `body_hash_index` is compared to the the total field sequence as seen below:

```

1 assert(
2     body_hash_index > dkim_header_field_sequence.index
3     & body_hash_index < dkim_header_field_sequence.end_index(), "Body hash index
4     accessed outside of DKIM header field"
5 );

```

**Snippet 4.7:** Snippet from `lib/src/headers/body_hash.nr:get_body_hash()`

However, the second comparison should instead compare with `body_hash_index + BODY_HASH_BASE64_LENGTH`. The current comparison only ensures the start of the read is within the bounds, instead of the entire read existing within the bounds.

**Impact** Although one could not feasibly exploit the DKIM verification given the nature of the body hash field, users could insert the `bh=` string at an arbitrary point in the body hash in order to read past the intended field body. This may be an issues for consumers of the `zkemail.nr` library that want to stipulate about the body hash field of a header.

**Recommendation** Change the second constraint to check `body_hash_index + BODY_HASH_BASE64_LENGTH < dkim_header_field_sequence.end_index()`.

**Developer Response** The recommended fix was implemented so that the body hash is constrained to be contained within the DKIM signature header field.

#### 4.1.12 V-ZIM-VUL-012: Email nullifiers may leak information on the message

|                         |                     |               |                    |
|-------------------------|---------------------|---------------|--------------------|
| <b>Severity</b>         | Warning             | <b>Commit</b> | 2f81196            |
| <b>Type</b>             | Information Leakage | <b>Status</b> | Acknowledged       |
| <b>File(s)</b>          |                     |               | lib/src/lib.nr     |
| <b>Location(s)</b>      |                     |               | standard_outputs() |
| <b>Confirmed Fix At</b> |                     |               | N/A                |

The `standard_outputs()` function will output the hash of the pubkey and a hash of the signature, in which the latter is intended to be used as a nullifier:

```

1 pub fn standard_outputs<let KEY_BYTE_LENGTH: u32>(
2     pubkey: [Field; KEY_BYTE_LENGTH],
3     signature: [Field; KEY_BYTE_LENGTH]
4 ) -> [Field; 2] {
5     let pubkey_hash = pedersen_hash(pubkey);
6     let email_nullifier = pedersen_hash(signature);
7     [pubkey_hash, email_nullifier]
8 }

```

**Snippet 4.8:** Snippet from `lib/src/lib.nr:standard_outputs()` with comments removed

Given that RSA signatures are not hiding, an information leakage may occur if the signature is leaked. If an attacker is able to find which public key signed the message (which can be done by attempting to verify the signature against a database of DKIM public keys), then the note emitted is known to have been related to this signing entity. Additionally, RSA signatures are deterministic and dependent upon the contents being signed. Therefore, if the attacker is aware of the structure of a message (such as in `z-imburse`), and they know a list of potential recipients, they could theoretically brute-force the contents of the message by finding the matching signature.

**Impact** Information leakage on the email which is nullified may occur if the signature is leaked out of band. Note that this function is not currently used anywhere except in examples.

Additionally, emails may commonly be resent (e.g. requesting a new copy of a receipt), which will likely cause the signature to change due to the new email timestamp. Consequently, the nullifier may not be effective for its intended purpose of preventing email replays.

**Recommendation** Add a deterministic nonce to the `email_nullifier` which is not dependent upon the pubkey or signature.

**Developer Response** Since the email nullifier is hashed, it should not be possible for an attacker to determine the pre-image given the nullifier. In the context of Aztec contracts, DKIM public key hashes are accessed in the private context and observers have no insight into the public key used. Additionally, the signature will sign the exact datetime of transmission on top of the variance in the body, meaning we believe the domain for brute forcing the pre-image is sufficiently large enough that this attack is likely not feasible. We are of course readily eager to be challenged on this assertion, but this is our view to the best of our knowledge currently.

Additionally, ZKEmail in Circom uses  $H(H(\text{signature}))$  to create an email nullifier. While we only use one round of hashing and may need to use two, our goal is to emulate as closely as possible the decisions and features of ZKEmail in circom.

Re: replayed emails, this issue is separate from the function of this email nullifier type. This is a separate and potentially serious issue that warrants its own concerns but is not meant to be mitigated by the email nullifier. The email nullifier is simply supposed to prevent the exact used email from being used twice - not a clone sent at a different time. If the auditor contests our first opinion, we will consult with the ZKEmail team on this issue. Otherwise, we do not plan to take any action on this ticket.

**Updated Veridise Response** The issue is slightly more nuanced than relying on the preimage resistance of the hash function. Also note that as we mentioned, the `standard_outputs()` isn't actually used anywhere in Z-imburse except for the examples.

The issue is that we are working under the assumption that somehow the signature *does* get leaked out of band, whereas the contents of the email remain private. We acknowledge this is an extremely unlikely scenario to occur.

However, if this does happen (and the `standard_outputs()` function is used in Z-imburse) then when the reimbursement transaction is executed and the nullifiers are emitted, an attacker can deduce which provider the reimbursement was for. Additionally, given the strict structure of the emails being verified, an attacker may be able to brute force the contents of the email and attempt to verify the contents against the public key + signature until signature verification passes. At that point he has uncovered the contents of the email and therefore who was reimbursed, from which provider, and for how much.

#### 4.1.13 V-ZIM-VUL-013: Ignored DKIM Tags

|                         |                 |               |                       |
|-------------------------|-----------------|---------------|-----------------------|
| <b>Severity</b>         | Warning         | <b>Commit</b> | 2f81196               |
| <b>Type</b>             | Data Validation | <b>Status</b> | Acknowledged          |
| <b>File(s)</b>          |                 |               | See issue description |
| <b>Location(s)</b>      |                 |               | See issue description |
| <b>Confirmed Fix At</b> |                 |               | N/A                   |

Several DKIM tags (see [Sec 3.5 of RFC 6376](#)) are ignored, with varying consequences.

1. The canonicalization tag `c=` is ignored. It is assumed that a relaxed canonicalization is used for the header.
2. The version tag `v=` is not checked for validity.
3. The body length tag `l=` is not used.
4. The expiration tag `x=` is ignored.
5. The SDID tag `d=` is ignored. Note that RFC 6376 dictates that applications are required to check this tag to distinguish between multiple entities using the same DKIM key.

**Impact** Assuming relaxed canonicalization or ignoring the `l=` tag may lead to denial-of-service issues.

Accepting unexpected versions or expired DKIM signatures may cause applications to accept invalid emails.

**Recommendation** Add documentation describing the ignored tags, and what additional checks a library user may need to perform.

**Developer Response** The developers have been notified of the issue but have yet to respond with acknowledgement or fixes.

#### 4.1.14 V-ZIM-VUL-014: Optimization opportunities

|                         |   |               |         |
|-------------------------|---|---------------|---------|
| <b>Severity</b>         | Info  | <b>Commit</b> | 2f81196 |
| <b>Type</b>             | Optimization  | <b>Status</b> | Fixed   |
| <b>File(s)</b>          | See issue description   |               |         |
| <b>Location(s)</b>      | See issue description   |               |         |
| <b>Confirmed Fix At</b> | <a href="https://github.com/zkemail/zkemail.nr/pull/31">https://github.com/zkemail/zkemail.nr/pull/31</a> |               |         |

In the following locations, the auditors identified missed opportunities for constraint optimizations:

- ▶ `lib/src/remove_soft_line_breaks.nr`:
  - `find_zeroes()`: Since `is_break[N - 2]` will always be zero, the expression `should_zero[N - 1] = is_break[N - 2] + is_break[N - 3]`; can be simplified to `should_zero[N - 1] = is_break[N - 3]`;

**Impact** Circuits may perform poorly or gas may be wasted, costing users extra time and funds.

**Recommendation** Perform the optimizations.

**Developer Response** We implemented the recommendation.



## Glossary

**DKIM** DomainKeys Identified Mail (DKIM) Signatures are commonly used to authenticate e-mail senders. See IETF RFC 6376 to learn more (<https://datatracker.ietf.org/doc/html/rfc6376>). 1

**Noir** A DSL for writing zero-knowledge circuits in a high-level, rust-like syntax. See <https://noir-lang.org> to learn more. 1

**RSA** A digital signature scheme. See <https://datatracker.ietf.org/doc/html/rfc3447> to learn more. 1

**zero-knowledge circuit** A cryptographic construct that allows a prover to demonstrate to a verifier that a certain statement is true, without revealing any specific information about the statement itself. See [https://en.wikipedia.org/wiki/Zero-knowledge\\_proof](https://en.wikipedia.org/wiki/Zero-knowledge_proof) for more. 1