Security Assessment of
Relaycorp's Relaynet Network Protocol
(on behalf of the Open Technology Fund)

Relaycorp

OPEN TECHNOLOGY FUND

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Scope and Methodology

IncludeSec performed a security assessment of Relaycorp's Relaynet Network Protocol (on behalf of the Open Technology Fund). The assessment team performed a 5 day effort spanning from March 18th – March 22nd, 2019, using a Time-Boxed Grey Box Assessment Methodology which included a detailed review of all the components described above in a manner consistent with the original Statement of Work (SOW).

Assessment Objectives

The objective of this assessment was to identify and confirm potential security vulnerabilities within targets in-scope of the SOW. The team assigned a qualitative risk ranking to each finding. IncludeSec also provided remediation steps which Relaycorp could implement to secure its applications and systems.

Findings Overview

IncludeSec identified 7 categories of findings. There were 0 deemed a "Critical-Risk," 4 deemed a "High-Risk," 1 deemed a "Medium-Risk," and 1 deemed a "Low-Risk," which pose some tangible security risk. Additionally, 1 "Informational" level findings were identified that do not immediately pose a security risk.

IncludeSec encourages Relaycorp to redefine the stated risk categorizations internally in a manner that incorporates internal knowledge regarding business model, customer risk, and mitigation environmental factors.

Next Steps

IncludeSec advises Relaycorp to remediate as many findings as possible in a prioritized manner and make systemic changes to the Software Development Life Cycle (SDLC) to prevent further vulnerabilities from being introduced into future release cycles. This report can be used by Relaycorp as a basis for any SDLC changes. IncludeSec welcomes the opportunity to assist Relaycorp in improving their SDLC in future engagements by providing security assessments of additional products.

ASSESSMENT RESULTS

At the conclusion of the assessment, Include Security categorized findings into four levels of perceived security risk: critical, high, medium, or low. Any informational findings for which the assessment team perceived no direct security risk, were also reported in the spirit of full disclosure. The risk categorizations below are guidelines that IncludeSec believes reflect best practices in the security industry and may differ from internal perceived risk. It is common and encouraged that all clients recategorize findings based on their internal business risk tolerances. All findings are described in detail within the final report provided to Relaycorp.

Critical-Risk findings are those that pose an immediate and serious threat to the company's infrastructure and customers. This includes loss of system, access, or application control, compromise of administrative accounts or restriction of system functions, or the exposure of confidential information. These threats should take priority during remediation efforts.

High-Risk findings are those that could pose serious threats including loss of system, access, or application control, compromise of administrative accounts or restriction of system functions, or the exposure of confidential information.

Medium-Risk findings are those that could potentially be used with other techniques to compromise accounts, data, or performance.

Low-Risk findings pose limited exposure to compromise or loss of data, and are typically attributed to configuration issues, and outdated patches or policies.

Informational findings pose little to no security exposure to compromise or loss of data which cover defense-in-depth and best-practice changes which we recommend are made to the application.

The findings below are listed by a risk rated short name (e.g., C1, H2, M3, L4, I5) and finding title. Each finding includes: Description (including proof of concept screenshots and lines of code), Recommended Remediation, and References.

Project Scoping, Threat Modeling, and Assessment Methodology

Project Scoping

On March 18th, 2019, the assessment team began analyzing Relaynet for security vulnerabilities (see reference links below for exact versions.). The assessment was time boxed to 48 hours. Relaynet's specification was solely in scope the Relaynet implementation received by the IncludeSec team was only to be used to verify findings found within the specification.

Given that the time spent on the assessment was not exhaustive, IncludeSec hopes this project serves as inspiration for the open source communities to execute their own security reviews and to report any identified vulnerabilities to the project maintainers.

Threat Modeling

As Relaynet is meant to define a secure messaging protocol which can operate in an untrusted environment, threat modeling included actors who can partially or fully partake in generating, receiving, and transferring messages. Note that while Relaynet makes heavy use of public key cryptography, initial exchange of the Gateways's public key was out of scope of the assessment. The following areas were of key focus during the assessment:

- Backdoors – Assessing if the specification voluntary or involuntary contains backdoors, which allow for spoofing and tampering of messages, recovery of plaintext and keys, and determining message similarity based on ciphertext. Examples include weak parameters, use of broken cryptographic algorithms, and key mismanagement.

- Correctness – Assessing if the specification is detailed and concise. Statements which are ambiguous or contradicting may lead to security vulnerabilities when implemented.

Testing Methodology

As Relaynet is meant to define a secure messaging protocol which can operate in an untrusted environment, prior research regarding vulnerabilities found within Secure Socket Layer and Transport Layer Security specifications were reviewed. Additionally, security considerations for Signal's Double Ratchet Algorithm and Extended Triple Diffie-Hellman specification were reviewed, as Relaynet uses the specifications with the goal to obtain perfect forward secrecy, future secrecy and replay attack mitigation. Relaynet findings were verified against the supplied implementation when possible. Appropriate proofs-of-concept were developed to verify discovered findings. Please also note that the level of depth of attacks was limited by the time-boxed nature of the assessment (6 workdays in total.)

References

Relaynet Reference Specification
Summarizing Known Attacks on Transport Layer Security (TLS) and Datagram TLS (DTLS)
The Double Ratchet Algorithm
The X3DH Key Agreement Protocol

CRITICAL-RISK FINDINGS

No critical-risk findings identified in the project time window within the defined scope.

HIGH-RISK FINDINGS

H1: Relaynet Abstract Message Format Hashing Algorithm Not Specified

Description:

The Relaynet Abstract Message Format specification (**rs001-ramf**) describes a signature field used to verify the contents of a message. Acceptable hashing algorithms are not listed within the specification, which may lead to deprecated hashing algorithms being used, such as **MD5** or **none**.

The following is the quoted specification relevant to the finding, which is from https://github.com/relaynet/specs/blob/a7f7b8ced225950cf2e1c9552f383a5124b06238/rs001-ramf.md, line 24:

```
2. Signature hashing algorithm, defined early to allow the recipient to start calculating the
message digest as the message is being streamed. This is an ASCII with a fixed length of 8
octets, padded with 0x00 octets at the end if fewer octets are needed.
```

Recommended Remediation:

The assessment team recommends specifying acceptable hash algorithms which can be used based on security considerations. The hashing algorithms **SHA-256** or **SHA-512** are commonly used, depending on security needs.

References:

CWE-327: Use of a Broken or Risky Cryptographic Algorithm

H2: Minimum Required Key Strength Not Specified

Description:

The Relaynet Channel Session Protocol specification (rs003-key-agreement) states RSA can be used during the key agreement protocol. No acceptable RSA key sizes are defined, allowing keys less than 2048 bits to be used. Additionally, the specification states a curve like Curve25519 or Curve448 can be used. The phrase a curve like is ambiguous. Accepted curves should be explicitly listed to ensure a weak curve can't be selected.

The following is the quoted specification relevant to the finding, which is from https://github.com/relaynet/specs/blob/a7f7b8ced225950cf2e1c9552f383a5124b06238/rs003-key-agreement.md, line 33:

```
Algorithm: RSA or a curve like Curve25519 or Curve448
```

Recommended Remediation:

The assessment team recommends explicitly stating acceptable key strengths to protect communications. At a minimum, keys used by asymmetric cryptography should provide an equivalent key strength of 128-bit symmetric keys. Symmetric keys should provide at least 128-bit security.

References:

CWE-326: Inadequate Encryption Strength

H3: Diffie-Hellman Key Exchange may Operate on Key Material Generated from Different Algorithms and Parameters

Description:

The Relaynet Channel Session Protocol specification (rs003-key-agreement) uses **Diffie-Hellman** for the key exchange. In the Key Agreement Protocol section, the use of a Public Key stored in a certificate is used to establish initial communications. As RSA and Elliptic Curves materials may be used, implementations may accidentally use key material generated from different algorithms or parameters, which may lead to compromised communications due to generation of vulnerable keys. Furthermore, proofs regarding security guarantees of using RSA keys to derive Diffie-Hellman keys were not found when researched.

The following is the quoted specification relevant to the finding, which is from https://github.com/relaynet/specs/blob/a7f7b8ced225950cf2e1c9552f383a5124b06238/rs003-key-agreement.md, on lines 80 and 96:

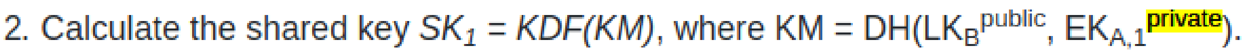

2. Calculate the shared key $SK_1 = KDF(KM)$, where KM = DH(LK_B^{public}, $EK_{A,1}^{private}$).

LK is the key pair associated with the public key within the certificate.

Recommended Remediation:

The assessment team recommends avoiding derived Diffie-Hellman keys from other keys. Instead, consider the use of static Diffie-Hellman parameters embedded in the certificate, which are signed and can be used for the initial key exchange. Additionally, Diffie-Hellman key exchanges using newly generated ephemeral keys should occur for forward secrecy, which is already implemented as part of the spec under Sending Subsequent Messages and Receiving Subsequent Messages. Additionally, ensure cryptographic operations use the same groups and parameters, wherever required for security guarantees.

References:

CWE-327: Use of a Broken or Risky Cryptographic Algorithm

H4: Supported Symmetric Encryption Algorithms Not Specified

Description:

The Relaynet Channel Session Protocol specification (rs003-key-agreement) states encryption algorithms are specified by their **OID**. No Relaynet specification states which encryption algorithms must be used. For instance, a vulnerable cipher such as **DES**, which is assigned the **OID 1.3.36.3.1.1**, may be used to encrypt material. Note that **DES** have an effective size of 56 bits, which is brute forceable and may allow disclosure of cleartext data.

The following is the quoted specification relevant to the finding, which is from https://github.com/relaynet/specs/blob/a7f7b8ced225950cf2e1c9552f383a5124b06238/rs003-key-agreement.md, line 134:

```
keyEncryptionAlgorithm: The algorithm identified by its OID. The ASN.1 representation of this
identifier is shown below:
```

Recommended Remediation:

The assessment team recommends limiting acceptable ciphers based on security requirements. For encrypting communications, consider using a library which is easy to use and contains reasonable ciphers and parameters, such as **libsodium**.

References:

CWE-327: Use of a Broken or Risky Cryptographic Algorithm
Libsodium

MEDIUM-RISK FINDINGS

M1: Use of Public Hostnames Without Warning Message

Description:

The assessment team found that the documentation does not warn about the dangers of resolving hostnames using public DNS services.

According to https://github.com/relaynet/specs/blob/master/rs004-cosocket.md#binding-hint, the **cosocket** URL is defined as the following:

- The hint for this binding MUST be **cosocket**. For example, **rng+cosocket://example.com** would be a valid public gateway address.

Which is defined here: https://github.com/relaynet/specs/blob/master/rs004-cosocket.md

```
Gateway Messaging Protocol
This protocol establishes the channel between two gateways.
Gateway addresses MUST use the scheme rng. For example, rng://example.com and
rng+grpc://example.com (if using the gRPC binding) are valid public gateway addresses, and
rng:0b5bb9d8014a0f9b1d61e21e796d78dccdf1352f23cd32812f4850b878ae4944c is a valid private
gateway address.
When using the Relaynet Key Agreement protocol, the two gateways MUST maintain a single
session across the different message types.
```

The specification does not add any details whether a public gateway address is assumed to be used in a safe environment or not.

It's possible that external malicious third-party DNS servers would be queried and therefore would be able to identify a Cargo exchange. Furthermore, it would be possible to respond with a malicious IP that could take advantage of a connection with the relay client.

Recommended Remediation:

In a sensitive and complex situation, the assessment team recommends specifying that DNS servers must be trusted or within a trusted environment.

LOW-RISK FINDINGS

L1: Abstract Message Format Signature Check Process Not Explicitly Defined

Description:

The Relaynet Abstract Message Format specification (**rs001-ramf**) uses the phrase Check the signature to describe the process of verifying the integrity of a message. The phrase may be interpreted in many ways, such as ensuring the field is not **null**.

The following is the quoted specification relevant to the finding, which is from https://github.com/relaynet/specs/blob/a7f7b8ced225950cf2e1c9552f383a5124b06238/rs001-ramf.md, line 46:

```
Check the signature.
```

Recommended Remediation:

The assessment team recommends expanding on what the phrase Check the signature entails.

References:

CWE-347: Improper Verification of Cryptographic Signature

INFORMATIONAL FINDINGS

I1: Missing Threats Diagram

Description:

Before starting the security assessment based on the current specifications, the assessment team performed a light threat analysis based on the documents shared by Relaynet for secure design review.

The assessment team defined a list of points of interest by locating a set of potentially vulnerable components that were taken into consideration during the diagrams and specs analysis. The following diagram indicates points of critical vulnerabilities where a security concern might be considered.

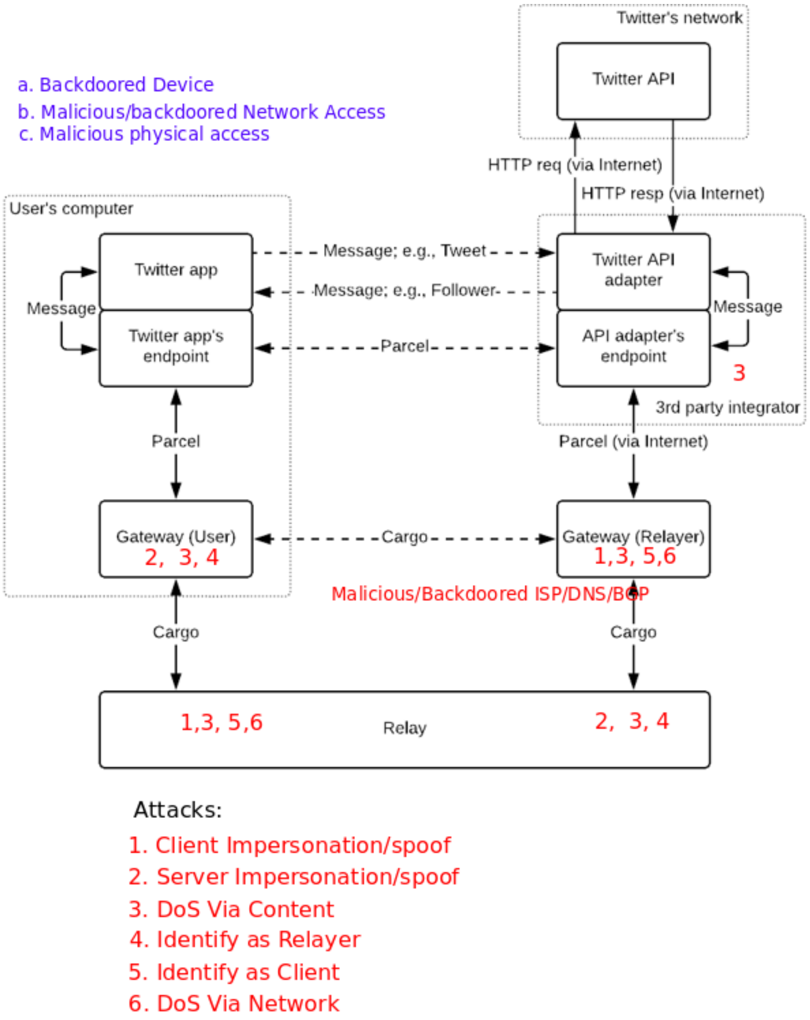

According to Relaynet documents, the diagram displays points of interest where the following factors are involved:

- User devices
- Gateway
- Relay device
- Relayer's Gateway

Considering the above factors, the following scenarios were identified (labeled in purple on the diagram):

a. Backdoored devices: Devices that can't be trusted because of some integrity loss.
b. Malicious/Backdoored Network: ISP, DNS, BGP: Networks that can't be trusted because of possible compromise or unethical activity.
c. Malicious physical access: An attacker has access to the targeted user's device to some extent.

Considering each one of the three previous scenarios, the diagram explains where it is recommended to create an abuse case table that can help identify critical threats.

Each scenario identifies a set of threats that can be summarized as follows (labeled in red on the diagram):

1. Client impersonation/spoof: In scenario a (backdoored device): If the device is compromised, a malicious user can impersonate the client and send malicious data acquire original data.
2. Server impersonation/spoof: In scenario b (malicious/backdoored network access): If the communication endpoint is compromised, a malicious server can acquire original data.
3. DoS via Content: In scenarios a (backdoored device), b (malicious/backdoored network access), and c (malicious physical): If a device or communication endpoints are compromised, the client or server should have monitoring services that detect a DoS attack might be ongoing. Additionally, a malicious third party with physical access to the device would be able to potentially execute actions that would result in a DoS.
4. Identify as Relayer: In scenarios a (backdoored device), b (malicious/backdoored network access), and c (malicious physical): The same vulnerabilities could be an issue as defined in 1 and 2.
5. Identify as Client: In scenarios a (backdoored device), b (malicious/backdoored network access), and c (malicious physical): The same vulnerabilities could be an issue as defined in 1 and 2.
6. DoS via Network: In scenarios a (backdoored device), b (malicious/backdoored network access), and c (malicious physical): The same vulnerabilities could be an issue as defined in 3.

Using the above diagram in Relaynet documentation, it would help the reader minimize implementation mistakes and help clarify unsuspected security situations.

Recommended Remediation:

The assessment team recommends securing design and implementation using the following specifications:

- Considering assets and factors, be sure to identify a well-defined set of controls that will have to be followed during implementation.
- Take into consideration applying to an RFC as it would be formally reviewed without missing important details.