RATS Working Group
Internet-Draft
Updates: 9334 (if approved)
Intended status: Informational
Expires: 12 August 2026

M. U. Sardar
TU Dresden

8 February 2026

Guidelines for Security Considerations of RATS
draft-sardar-rats-sec-cons-02

Abstract

This document aims to provide guidelines and best practices for writing security considerations for technical specifications for RATS targeting the needs of implementers, researchers, and protocol designers. This is a work-in-progress, and the current version mainly presents an outline of the topics that future versions will cover in more detail.

* Corrections in published RATS RFCs

* Security concerns in two RATS drafts

* General security guidelines, baseline, or template for RATS

About This Document

This note is to be removed before publishing as an RFC.

The latest revision of this draft can be found at https://muhammad-usama-sardar.github.io/rats-sec-cons/draft-sardar-rats-sec-cons.html. Status information for this document may be found at https://datatracker.ietf.org/doc/draft-sardar-rats-sec-cons/.

Source for this draft and an issue tracker can be found at https://github.com/muhammad-usama-sardar/rats-sec-cons.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

1.1. Need for Specialized Guidance in RATS

Every Internet Draft needs to have a "Security Considerations" section. While general guidelines such as [RFC3552] exist, the underlying threat model is that the endpoint is fully trusted (i.e., all software and hardware components in the device may access the keys). RATS [RFC9334] has a primarily different threat model in the sense that only parts of the endpoint (called Attester) are trusted (i.e., only specific software and hardware components in the device may access the keys), and the goal is to establish the trustworthiness of the endpoint. In other words, [RFC3552] deals with a network adversary, whereas RATS deals with an endpoint adversary, which may have root access or physical control over the device with which it can extract keys from software or hardware.

Moreover, remote attestation has several distinguishing features that necessitate a separate document. One specific example of such a feature is the architectural complexity of the endpoint. While network protocols typically have 2 roles, RATS has additional roles, which complicates the picture. Unfortunately, no guidelines currently exist for remote attestation [RFC9334] in RATS. This document aims to fill this gap.

1.2. Needs of the Target Audience of RATS

Moreover, while the target audience of Internet Drafts is implementers, researchers, and protocol designers [I-D.irtf-cfrg-cryptography-specification], RATS drafts generally do not fulfill these needs, in particular the needs of researchers and protocol designers. On the other hand, in our observation, implementers generally find it hard to relate the abstract concepts of RATS to the real-world systems. In general, implementers and protocol designers of RATS are thus left with little or no guidance.

1.3. Inaccuracies in Published RATS RFCs

Unfortunately, many published RFCs of RATS provide inaccurate or ambiguous security and privacy considerations, which may lead to errors in design and implementation, and give a false sense of security. As an example, many proposed designs in [RFC9334] are broken.

1.4. Aggregator in CoServ

RATS has recently adopted [I-D.ietf-rats-coserv], which has an ambiguous role Aggregator, for which -- in our assessment -- the authors have not yet provided a reasonable justification. To the best of our knowledge and understanding, a malicious Aggregator breaks the security of the RATS ecosystem and invalidates the formal proofs for RATS primitives. Surprisingly, during the three-week adoption call and one week discussion afterwards, one of the authors of the draft [I-D.ietf-rats-coserv] did not support adoption of the draft. Based on the above reasons, as researchers, we have genuine skepticism about this work. We request the authors to be transparent on this work and clarify the concerns raised at the adoption time (summarized to some extent in this draft).

1.5. Scope

To improve the situation, this draft presents an outline of three topics that future versions will cover in more detail:

* Corrections in published RATS RFCs [RFC9334], [RFC9781], [RFC9783] and [RFC9711]

* Security concerns in one currently adopted RATS draft [I-D.ietf-rats-coserv] and one proposed for adoption RATS draft [I-D.deshpande-rats-multi-verifier]

* General security baseline that other drafts can simply point to, or guidelines or template that other drafts can use

2. Conventions and Definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. General Hierarchy of Authentication

Authentication is a term which is often ambiguous in RATS specifications. We propose general hierarchy of one-way authentication [Gen-Approach], which can help precisely state the intended level of authentication (in decreasing order):

* One-way injective agreement
* One-way non-injective agreement
* Aliveness

Recentness can be added to each of these levels of authentication. Details will be added in future versions.

4. Threat Modeling

This section describes "What can go wrong?" TODO.

4.1. System Model

TODO.

4.2. Actors

TODO.

4.2.1. Legal perspective

* Data subject is an identifiable natural person (as defined in Article 4 (1) of GDPR [GDPR]).
* (Data) Controller (as defined in Article 4 (7) of GDPR [GDPR]) manages and controls what happens with personal data of data subject.
* (Data) Processor (as defined in Article 4 (8) of GDPR [GDPR]) performs data processing on behalf of the data controller.

TODO.

4.2.2. Technical perspective

* Infrastucture Provider is a role which refers to the Processor in GDPR. An example of this role is a cloud service provider (CSP).

TODO.

4.3. Threat Model

TODO.

4.4. Typical Security Goals

TODO.

5. Attacks

Security considerations in RATS specifications need to clarify how the following attacks are avoided or mitigated:

5.1. (Evidence) Replay Attacks

In this attack, a network or endpoint adversary -- with access to older Evidence -- can replay Evidence with stale Claims which no longer represent the actual state of the Attester, potentially resulting in exposure of confidential data [RA-TLS].

Replay of stale Evidence may be within the same connection or across multiple connections.

5.2. Diversion Attacks

In this attack, a network adversary -- with Dolev-Yao capabilities [Dolev-Yao] and access (e.g., via Foreshadow [Foreshadow]) to the attestation key of any machine in the world -- can redirect a connection intended for a specific Infrastructure Provider to the compromised machine, potentially resulting in exposure of confidential data [ID-Crisis].

In the context of confidential computing and TLS as a transport protocol, we reported these attacks to the TLS WG in February 2025 [Usama-TLS-26Feb25]. A formal proof is available [ID-Crisis-Repo] for further research and development. Since reporting to TLS WG, these attacks have been practically exploited in TEE.fail (https://tee.fail/), Wiretap.fail (https://wiretap.fail/), and BadRAM (https://badram.eu/).

5.3. Relay Attacks

In this attack, a network or endpoint adversary -- with access to suitable binding material -- can relay an attestation request to a genuine Attester and present the genuine Evidence as its own, potentially resulting in impersonation of genuine Attester [RelayAttacks-RATS].

Note that _replay_ is about _same_ Attester while _relay_ attack is about _different_ Attesters.

6. Potential Mitigations

This section describes the countermeasures and their evaluation.

To mitigate the above attacks, we propose post-handshake attestation. We are not aware of any attacks on post-handshake attestation. Post-handshake attestation avoids replay attacks by using a fresh attestation nonce. Moreover, considering TLS as the transport protocol, it avoids diversion and relay attacks by binding the Evidence to the underlying TLS connection, such as using Exported Keying Material (EKM) [I-D.ietf-tls-rfc8446bis], as proposed in Section 9.2 of [ID-Crisis]. [RFC9261] and [RFC9266] provide mechanisms for such bindings. Efforts for a formal proof of security of post-handshake attestation are ongoing.

7. Examples of Specifications That Could Be Improved

7.1. RFC9334

7.1.1. Unprotected Evidence

Section 7.4 of [RFC9334] has:

> A conveyance protocol that provides authentication and integrity protection can be used to convey Evidence that is otherwise unprotected (e.g., not signed).

Using a conveyance protocol that provides authentication and integrity protection, such as TLS 1.3 [RFC8446], to convey Evidence that is otherwise unprotected (e.g., not signed) undermines all security of remote attestation. Essentially, this breaks the chain up to the trust anchor (such as hardware manufacturer) for remote attestation. Hence, remote attestation effectively provides no protection in this case and the security guarantees are limited to those of the conveyance protocol only. In order to benefit from remote attestation, Evidence MUST be protected using dedicated keys chaining back to the trust anchor for remote attestation.

7.1.2. Missing definitions

[RFC9334] uses the term Conceptual Messages in capitalization without proper definition.

7.1.3. Missing Roles and Conceptual Messages

* Identity Supplier and its corresponding conceptual message Identity are missing and need to be added to the architecture [Tech-Concepts].

* Attestation Challenge as conceptual message needs to be added to the architecture [Tech-Concepts].

7.2. RFC9781

As argued above for RFC9334, security considerations in [RFC9781] are essentially insufficient.

7.3. RFC9783

[RFC9783] uses:

* 3x epoch handle (with reference to Section 10.2 of [RFC9334] and Section 10.3 of [RFC9334]) whereas RFC9334 never uses epoch handle at all!

* 1x epoch ID with no reference and no explanation of how it is different from epoch handle

7.4. RFC9711

7.4.1. Inaccurate opinion

Section 7.4 of [RFC9711] has:

> For attestation, the keys are associated with specific devices and are configured by device manufacturers.

The quoted text is inaccurate and just an opinion of the editors. It should preferably be removed from the RFC. For example, in SGX, the keys are not configured by the manufacturer alone. The platform owner can provide a random value called OWNER_EPOCH.

For technical details and proposed text, see [Clarifications-EAT].

7.4.2. Inaccurate Privacy Considerations

Section 8.4 of [RFC9711] has:

> The nonce claim is based on a value usually derived remotely (outside of the entity).

Attester-generated nonce does not provide any replay protection since the Attester can pre-generate an Evidence that might not reflect the actual system state, but a past one.

See the attack trace for Attester-generated nonce at [Sec-Cons-RATS].

For replay protection, nonce should _always_ be derived remotely (for example, by the Relying Party).

8. Examples of Parts of Specifications That are Detrimental for Security

We believe that the following parts of designs are detrimental for the RATS ecosystem:

8.1. Multi-Verifiers

8.1.1. Security Considerations

We believe the security considerations of multi-verifiers [I-D.deshpande-rats-multi-verifier] must say:

Compared to a single verifier, the use of multi-verifiers increases security risks in terms of increasing the Trusted Computing Base (TCB).

8.1.2. Privacy Considerations

We believe the privacy considerations of multi-verifiers [I-D.deshpande-rats-multi-verifier] should say:

Compared to a single verifier, the use of multi-verifiers may increase the privacy risks, as potentially sensitive information may be sent to multiple verifiers.

8.1.3. Open-source

Besides, the rationale presented by the authors at meeting 124 -- appraisal policy being the intellectual property of the vendors -- breaks the open-source nature of RATS ecosystem. This requires blindly trusting the vendors and increases the attack surface.

8.2. Aggregator-based design

Aggregator in [I-D.ietf-rats-coserv] is an explicit trust anchor and the addition of new trust anchor needs to have a strong justification. Having a malicious Aggregator in the design trivially breaks all the guarantees. It should be clarified how trust is established between Aggregator and Verifier in the context of Confidential Computing threat model.

The fact that Aggregator has collective information of Reference Values Providers and Endorsers makes it a special target of attack, and thus a single point of failure. It increases security risks because Aggregator can be compromised independent of the Reference Values Providers and Endorsers. That is, even if Reference Values Providers and Endorsers are secure, the compromise of Aggregator breaks the security of the system. Moreover, if Aggregator is not running inside a TEE, it is relatively easy to compromise the secrets.

9. Security Considerations

All of this document is about security considerations.

10. IANA Considerations

This document has no IANA actions.

11. References

11.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/rfc/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/rfc/rfc8174>.

[RFC9334] Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote ATtestation procedureS (RATS) Architecture", RFC 9334, DOI 10.17487/RFC9334, January 2023, <https://www.rfc-editor.org/rfc/rfc9334>.

[RFC9711] Lundblade, L., Mandyam, G., O'Donoghue, J., and C. Wallace, "The Entity Attestation Token (EAT)", RFC 9711, DOI 10.17487/RFC9711, April 2025, <https://www.rfc-editor.org/rfc/rfc9711>.

[RFC9781] Birkholz, H., O'Donoghue, J., Cam-Winget, N., and C. Bormann, "A Concise Binary Object Representation (CBOR) Tag for Unprotected CBOR Web Token Claims Sets (UCCS)", RFC 9781, DOI 10.17487/RFC9781, May 2025, <https://www.rfc-editor.org/rfc/rfc9781>.

[RFC9783] Tschofenig, H., Frost, S., Brossard, M., Shaw, A., and T. Fossati, "Arm's Platform Security Architecture (PSA) Attestation Token", RFC 9783, DOI 10.17487/RFC9783, June 2025, <https://www.rfc-editor.org/rfc/rfc9783>.

11.2. Informative References

[Clarifications-EAT]
Sardar, M. U., "Clarifications in draft-ietf-rats-eat", April 2025, <https://mailarchive.ietf.org/arch/msg/rats/4V2zZHhk5IuxwcUMNWpPBpnzpaM/>.

[Dolev-Yao]
Dolev, D. and A. Yao, "On the security of public key protocols", March 1983.

[Foreshadow]
Jo Van Bulck, Marina Minkin, Ofir Weisse, Daniel Genkin, Baris Kasikci, Frank Piessens, Mark Silberstein, Thomas F Wenisch, Yuval Yarom, and Raoul Strackx, "Foreshadow", October 2025, <https://foreshadowattack.eu/>.

[GDPR] European Commission, "Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Direc- tive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance)", May 2016, <https://eur-lex.europa.eu/eli/reg/2016/679/oj>.

[Gen-Approach]
Sardar, M. U., "Perspicuity of Attestation Mechanisms in Confidential Computing: General Approach", October 2025, <https://www.researchgate.net/publication/396593308_Perspicuity_of_Attestation_Mechanisms_in_Confidential_Computing_General_Approach>.

[I-D.deshpande-rats-multi-verifier]
Deshpande, Y., jun, Z., Labiod, H., and H. Birkholz, "Remote Attestation with Multiple Verifiers", Work in Progress, Internet-Draft, draft-deshpande-rats-multi-verifier-04, 7 February 2026, <https://datatracker.ietf.org/doc/html/draft-deshpande-rats-multi-verifier-04>.

[I-D.ietf-rats-coserv]
Howard, P., Fossati, T., Birkholz, H., Kamal, S., Mandyam, G., and D. Ma, "Concise Selector for Endorsements and Reference Values", Work in Progress, Internet-Draft, draft-ietf-rats-coserv-02, 20 October 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-rats-coserv-02>.

[I-D.ietf-tls-rfc8446bis]
Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", Work in Progress, Internet-Draft, draft-ietf-tls-rfc8446bis-14, 13 September 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-tls-rfc8446bis-14>.

[I-D.irtf-cfrg-cryptography-specification]
Sullivan, N. and C. A. Wood, "Guidelines for Writing Cryptography Specifications", Work in Progress, Internet-Draft, draft-irtf-cfrg-cryptography-specification-02, 7 July 2025, <https://datatracker.ietf.org/doc/html/draft-irtf-cfrg-cryptography-specification-02>.

[ID-Crisis]
Sardar, M. U., Moustafa, M., and T. Aura, "Identity Crisis in Confidential Computing: Formal Analysis of Attested TLS", November 2025, <https://www.researchgate.net/publication/398839141_Identity_Crisis_in_Confidential_Computing_Formal_Analysis_of_Attested_TLS>.

[ID-Crisis-Repo]
Muhammad Usama Sardar, "Identity Crisis in Confidential Computing: Formal analysis of attested TLS protocols", <https://github.com/CCC-Attestation/formal-spec-id-crisis>.

[RA-TLS] Sardar, M. U., Niemi, A., Tschofenig, H., and T. Fossati, "Towards Validation of TLS 1.3 Formal Model and Vulnerabilities in Intel's RA-TLS Protocol", November 2024, <https://www.researchgate.net/publication/385384309_Towards_Validation_of_TLS_13_Formal_Model_and_Vulnerabilities_in_Intel's_RA-TLS_Protocol>.

[RelayAttacks-RATS] Sardar, M. U., "Relay Attacks in Intra-handshake Attestation for Confidential Agentic AI Systems", January 2026, <https://mailarchive.ietf.org/arch/msg/rats/6gbqx0XY8WYrH3Mx4vO8n2-uKgY/>.

[RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, DOI 10.17487/RFC3552, July 2003, <https://www.rfc-editor.org/rfc/rfc3552>.

[RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <https://www.rfc-editor.org/rfc/rfc8446>.

[RFC9261] Sullivan, N., "Exported Authenticators in TLS", RFC 9261, DOI 10.17487/RFC9261, July 2022, <https://www.rfc-editor.org/rfc/rfc9261>.

[RFC9266] Whited, S., "Channel Bindings for TLS 1.3", RFC 9266, DOI 10.17487/RFC9266, July 2022, <https://www.rfc-editor.org/rfc/rfc9266>.

[Sec-Cons-RATS] Sardar, M. U., "Security considerations of remote attestation (RFC9334)", November 2024, <https://mailarchive.ietf.org/arch/msg/rats/jcAv9FKbYSIVtUNQ8ggEHL8lrmM/>.

[Tech-Concepts] Sardar, M. U., "Perspicuity of Attestation Mechanisms in Confidential Computing: Technical Concepts", October 2025, <https://www.researchgate.net/publication/396199290_Perspicuity_of_Attestation_Mechanisms_in_Confidential_Computing_Technical_Concepts>.

[Usama-TLS-26Feb25]
Muhammad Usama Sardar, "Impersonation attacks on protocol in draft-fossati-tls-attestation (Identity crisis in Attested TLS) for Confidential Computing", February 2025, <https://mailarchive.ietf.org/arch/msg/tls/Jx_yPoYWMIKaqXmPsytKZBDq23o/>.

Acknowledgments

The author wishes to thank Ira McDonald and Ivan Gudymenko for insightful discussions. The author also wishes to thank the authors of [I-D.ietf-rats-coserv] (in particular Thomas Fossati and Paul Howard) for several discussions, which unfortunately could not resolve the above concerns, and hence led to this draft.

History

-01

* Concrete text proposal for security and privacy considerations of multi-verifiers [I-D.deshpande-rats-multi-verifier]

-02

* Introduction and motivation

* Defined replay and relay attacks

* Added mitigations

Author’s Address

Muhammad Usama Sardar
TU Dresden
Email: muhammad_usama.sardar@tu-dresden.de