An Empirical Analysis of RSA Common-Factor Vulnerabilities in Contemporary TLS Certificates from Latin America

Análisis Empírico de Vulnerabilidades RSA por Factor Común en Certificados TLS de América Latina

DOI: http://dx.doi.org/10.17981/cesta.07.01.2026.05
Scientific research article
Date received: May 06, 2026, Date accepted: June 03, 2026.

Leonardo Lizcano Pinto
Universidad del Norte (Colombia)
ldlizcano@uninorte.edu.co
Daniela Ospino-Balcázar
Universidad del Norte (Colombia)
dospinoa@uninorte.edu.co

How to cite:
L. Vizcaino Pinto & D. Ospino Balcázar, "An Empirical Analysis of RSA Common-Factor Vulnerabilities in Contemporary TLS Certificates from Latin America", J. Comput. Electron. Sci.: Theory Appl., vol. 7, no. 1, pp. 53-60, 2026. DOI: 10.17981/cesta.07.01.2026.05

Abstract

Introduction: The security of RSA cryptographic systems critically depends on the proper generation of large prime numbers. Entropy failures during this process may cause different keys to share common prime factors, compromising private keys and exposing digital systems to cryptographic attacks.
Objective: To analyze whether the vulnerability associated with shared prime factors in RSA keys, reported in previous studies, persists in a contemporary context within TLS certificates linked to Latin American domains.
Method: A quantitative and computational approach was adopted, based on the cryptographic analysis of RSA public keys. A dataset of 1,444 valid RSA moduli was collected from TLS certificates associated with Latin American domains, using Certificate Transparency logs obtained through the crt.sh platform. Subsequently, an exhaustive pairwise greatest common divisor, GCD, computation was performed to identify possible shared prime factors among the analyzed moduli.
Results: The results showed no evidence of RSA moduli sharing non-trivial common factors within the sample analyzed. This indicates that the specific vulnerability reported in earlier studies, related to entropy failures and accidental reuse of prime factors, was not observed in the studied certificate set.
Conclusions: The findings suggest significant improvements in cryptographic key generation practices over the past decade, at least within the sample analyzed. However, the study highlights the importance of continuous cryptographic auditing to promptly detect potential weaknesses in digital security infrastructures.

Keywords: RSA; Cryptography; Common factors; TLS certificates; Certificate Transparency; GCD analysis

Resumen

Introducción: La seguridad de los sistemas criptográficos RSA depende de manera crítica de la generación adecuada de números primos grandes. Fallas en la entropía durante este proceso pueden provocar que distintas claves compartan factores primos comunes, comprometiendo la seguridad de las claves privadas y exponiendo sistemas completos a ataques criptográficos.
Objetivo: Analizar si la vulnerabilidad asociada al uso compartido de factores primos en claves RSA, reportada en estudios previos, persiste en un contexto contemporáneo dentro de certificados TLS vinculados a dominios latinoamericanos.
Método: Se adoptó un enfoque cuantitativo y computacional basado en el análisis criptográfico de claves públicas RSA. Para ello, se recolectó un conjunto de 1.444 módulos RSA válidos extraídos de certificados TLS asociados con dominios latinoamericanos, utilizando registros de Certificate Transparency obtenidos mediante la plataforma crt.sh. Posteriormente, se aplicó un cálculo exhaustivo por pares del máximo común divisor, GCD, con el fin de identificar posibles factores primos compartidos entre los módulos analizados.
Resultados: Los resultados no evidenciaron la existencia de módulos RSA que compartieran factores comunes no triviales dentro de la muestra analizada. Esto indica que la vulnerabilidad específica reportada en investigaciones anteriores, relacionada con fallas de entropía y reutilización accidental de factores primos, no se observó en el conjunto de certificados estudiado.
Conclusiones: Los hallazgos sugieren mejoras significativas en las prácticas de generación de claves criptográficas durante la última década, al menos en la muestra analizada. No

obstante, el estudio resalta la importancia de mantener procesos continuos de auditoría criptográfica para detectar oportunamente posibles debilidades en infraestructuras de seguridad digital.

Palabras clave

RSA; Criptografía; Factores comunes; Certificados TLS; Certificate Transparency; Análisis mediante GCD

INTRODUCTION

Decision-making in organizational contexts has undergone significant transformation over recent decades. The growing use of data, analytical models, and computational tools has shifted exclusively intuitive approaches toward more structured, evidence-based processes. Within this framework, artificial intelligence (AI) has emerged as a central component of organizational management, with applications ranging from strategic planning to day-to-day operations.

Public-key cryptography remains a foundational component of modern secure communication systems. Among these, RSA continues to be widely deployed in TLS/SSL infrastructure despite increasing adoption of elliptic-curve cryptography. The security of RSA relies fundamentally on the difficulty of factoring large composite integers generated as the product of two large, randomly chosen prime numbers.

In a seminal study, Heninger et al. [1] revealed that improper randomness during RSA key generation led to catastrophic failures, where millions of public keys shared common prime factors, enabling efficient factorization through simple greatest common divisor (GCD) computations. These vulnerabilities were particularly prevalent in embedded and headless devices suffering from insufficient entropy during boot-time key generation.

More than a decade has passed since these findings were published. Advances in operating systems, cryptographic libraries, hardware random number generators, and certificate authority policies may have mitigated such vulnerabilities. However, empirical evidence is required to assess whether these weaknesses persist in contemporary crypto-graphic infrastructures.

This paper investigates whether RSA keys with shared prime factors are still present in modern TLS certificates associated with Latin American domains. By replicating and adapting the methodology proposed by Heninger et al. [1], this work provides an updated regional assessment of RSA key robustness.

RELATED WORKS

The field of artificial intelligence applied to decision systems has evolved from approaches centered on traditional analytical models toward complex architectures based on machine learning, advanced analytics, and intelligent automation. This evolution reflects the convergence of multiple domains, including sustainable innovation, digital transformation, smart operations, and data-driven governance (Jordan & Mitchell, 2015; Ivanov et al., 2021; Huang & Rust, 2024).

From this perspective, knowledge analysis requires a systemic approach that integrates different thematic areas. The conceptual structure of the scientific literature can be understood through two complementary dimensions: (i) the thematic development areas of specialized journals and (ii) the emerging clusters derived from global bibliometric analysis (Braun et al., 2024; Shrestha et al., 2019; Sumbal et al., 2024).

The security of RSA-based public-key infrastructures critically depends on the availability of sufficient entropy during key generation. One of the most influential empirical studies in this area was conducted by Heninger et al. [1] who demonstrated that large-scale entropy failures in networked and embedded devices resulted in widespread RSA vulnerabilities. By performing Internet-wide scans of TLS and SSH servers, the authors showed that many RSA public keys shared common prime factors, enabling efficient factorization through greatest common divisor (GCD) computations and large-scale compromise of private keys.

Following this foundational work, subsequent studies investigated whether such vulnerabilities persisted after public disclosure. Yilek et al. [2] analyzed the aftermath of the Debian OpenSSL vulnerability, highlighting how weak entropy sources led to predictable RSA keys and how remediation efforts progressed unevenly across the ecosystem. These findings reinforced the idea that cryptographic weaknesses can persist long after they are identified, particularly in systems that are difficult to update or maintain.

Longitudinal analyses further explored the persistence of weak RSA keys in re-al-world deployments. Hastings et al. [3] conducted large-scale HTTPS measurements over multiple years and found that vulnerable cryptographic configurations, including weak RSA keys, often remained active due to slow patch adoption and vendor inertia. Their results emphasized that disclosure alone is insufficient to eliminate cryptographic weaknesses at scale.

More recent research suggests a gradual shift in the prevalence of RSA common-factor vulnerabilities, depending on the ecosystem under consideration. Kilgallin and Vasko [4] analyzed RSA keys across Internet-wide datasets and Certificate Transparency logs, concluding that shared-prime vulnerabilities had become increasingly rare in WebPKI certificates issued by major certificate authorities. However, they noted that such weaknesses remained more prevalent in constrained environments, particularly in IoT and embedded devices where entropy sources are limited. Similar conclusions were reported by Samara singhe and Mannan [5], who observed that modern TLS deployments benefit from stricter cryptographic policies and improved randomness generation, while legacy and embedded systems continue to pose residual risks.

In addition to common-factor vulnerabilities, other studies have examined related weaknesses in RSA deployments, such as key and certificate reuse. Nezhadian et al. [6] reported that RSA certificate and key reuse remains widespread across the public key infrastructure, even when keys are not directly factorable. Although reuse does not necessarily imply immediate cryptographic compromise, it may indicate operational weaknesses and increase systemic risk under certain threat models.

Taken together, the existing literature indicates that the large-scale RSA vulnerabilities reported in the early 2010s have been substantially mitigated in mainstream TLS eco-systems, largely due to improvements in entropy generation, cryptographic libraries, and certificate authority practices. Nevertheless, prior work consistently emphasizes the importance of continuous empirical measurement, as cryptographic weaknesses may persist or re-emerge in specific regions, device classes, or operational contexts. Within this re-search landscape, the present study contributes updated empirical evidence by evaluating RSA common-factor vulnerabilities in TLS certificates associated with Latin American domains using Certificate Transparency data.

MATERIAL AND METHODS

A. Methodological Workflow

The study adopts a quantitative bibliometric approach, aimed at analyzing the thematic structure and statistical relationship between the scientific production of IJMSOR and the Scopus database in the field of artificial intelligence applied to decision systems. The methodology adopted in this study followed a structured, multi-stage workflow designed to evaluate the presence of RSA common-factor vulnerabilities in contemporary TLS certificates. As illustrated in Figure 1, the process was organized into four sequential stages: certificate acquisition, dataframe construction, RSA modulus preparation, and cryptographic vulnerability analysis.

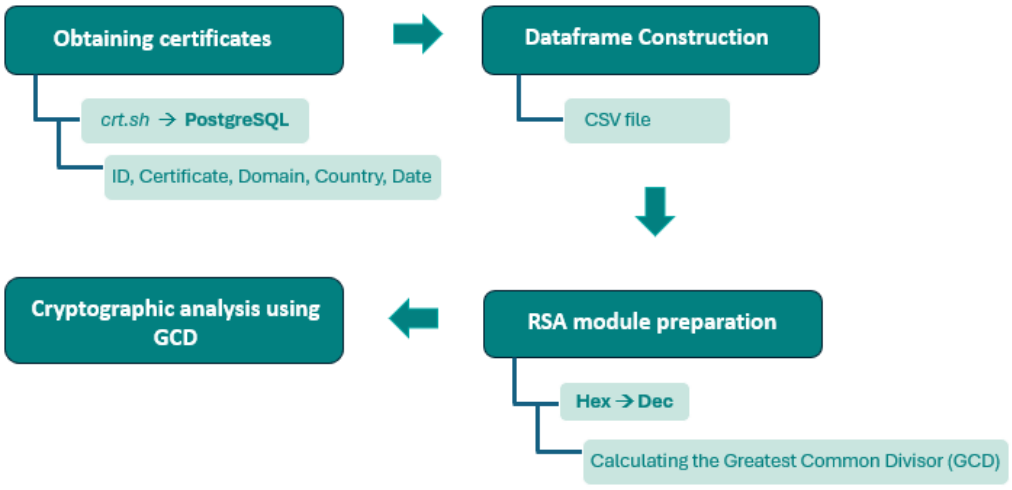

Figure 1. Methodological diagram.

At a high level, the workflow begins with the retrieval of publicly available TLS certificates from Certificate Transparency logs, followed by the construction of a curated dataset suitable for cryptographic analysis. The extracted RSA public keys are then prepared through numerical transformation and validation steps, enabling the application of arithmetic-based vulnerability detection techniques. Finally, a systematic cryptographic analysis is performed to identify potential shared prime factors among RSA moduli.

This workflow provides a reproducible and scalable framework for assessing the persistence of entropy-related RSA vulnerabilities in modern public-key infrastructures. De-tailed descriptions of each stage are provided in the following subsections.

B. Data Source

The dataset analyzed in this study was obtained from crt.sh, a public Certificate Transparency (CT) search engine that aggregates TLS certificates issued by certificate authorities worldwide. Certificate Transparency logs provide an auditable record of issued certificates and are widely used for security research.

A PostgreSQL query was executed on the crt.sh database to retrieve certificates issued within the last five years and associated with top-level domains (TLDs) corresponding to Latin American countries.

C. Dataset Construction

The initial query returned 1,499 TLS certificates, each containing the following information:

- Certificate identifier
- Full X.509 certificate
- Associated domain na
- Country or TLD
- Issuance date

Certificates not using RSA public keys or containing invalid or missing modulus values were excluded. After preprocessing, 55 certificates were removed due to null or malformed RSA parameters, resulting in a final dataset of 1,444 valid RSA public moduli.

D. RSA Modulus Preparation

RSA moduli extracted from X.509 certificates are typically encoded in hexadecimal format. Each modulus was converted into its corresponding integer representation to enable arithmetic operations required for cryptographic analysis.

E. GCD-Based Vulnerability Analysis

The core vulnerability analyzed in this study is based on the following principle: if two RSA modules share a prime factor, then they are vulnerable to complete factorization. Expressed mathematically:

$$\gcd(N_i, N_j) = p > 1 \Rightarrow N_i = p \cdot q_i,\ N_j = p \cdot q_j \quad (1)$$

An exhaustive pairwise GCD computation was performed across all RSA moduli in the dataset. Although more scalable batch-GCD algorithms exist, the moderate dataset size allowed a direct pairwise comparison without computational constraints. The complete analysis executed in approximately 36.8 seconds.

RESULTS

The GCD analysis of the 1,444 RSA moduli revealed no pairs sharing non-trivial common prime factors. For all evaluated pairs, the following condition held:

$$\gcd(N_i, N_j) = 1 \qquad \forall\, i \neq j \qquad (2)$$

Consequently, no RSA private keys could be derived using the common-factor attack described by Heninger et al [1].

In addition, only 48 certificates (3.32%) were found to reuse an identical RSA modulus. While key reuse may indicate configuration or deployment issues, it does not by itself imply vulnerability to factorization unless common primes are present.

The results obtained in this study differ markedly from those reported in the original work by Heninger et al. (2012), in which the authors identified widespread vulnerabilities affecting RSA keys at Internet scale. In that earlier study, a significant fraction of RSA public keys were found to share common prime factors, enabling efficient factorization and

large-scale compromise of private keys. By contrast, the analysis conducted on the present dataset did not reveal a single pair of RSA moduli sharing a non-trivial common factor.

Given the size and diversity of the analyzed sample, comprising TLS certificates associated with domains from multiple Latin American countries, it was reasonable to expect at least a small number of vulnerable keys exhibiting shared prime factors. However, the exhaustive pairwise GCD analysis performed over the 1,444 valid RSA moduli yielded no such coincidences. For all evaluated pairs, the greatest common divisor was equal to one, indicating that none of the RSA keys in the dataset were susceptible to the common-factor attack described in the original study.

This absence of vulnerable keys suggests that, for the period and regional scope considered, the specific RSA weakness reported in 2012 is no longer observable in contemporary TLS certificates. These findings are consistent with substantial improvements introduced over the past decade in cryptographic ecosystems. Such improvements include the widespread adoption of higher-quality hardware-based random number generators, enhanced entropy collection mechanisms in modern operating systems, stricter crypto-graphic policies enforced by certificate authorities, and improved key generation practices in widely used cryptographic libraries. Additionally, the gradual decommissioning of legacy and embedded devices known to suffer from entropy-related issues has likely contributed to the observed results.

A further notable contrast with the original study lies in the prevalence of repeated RSA keys. Heninger et al. reported that approximately 60% to 65% of RSA keys in their dataset were reused across multiple certificates, largely due to default keys or entropy failures in embedded systems. In the present analysis, only 48 certificates were found to share an identical RSA modulus with at least one other certificate, corresponding to approximately 3.3% of the dataset. While key reuse remains a potential concern from an operational or configuration standpoint, the significantly lower proportion observed here indicates a substantial reduction in this practice compared to the situation documented more than a decade ago.

Taken together, these results provide empirical evidence that the cryptographic infra-structure supporting modern TLS deployments has largely addressed the low-entropy and prime-reuse issues that previously undermined RSA security at scale. Nevertheless, the absence of detected vulnerabilities in this sample should not be interpreted as definitive proof that such issues have been entirely eradicated. Residual risks may persist in out dated systems, unmaintained devices, or constrained environments where entropy sources remain limited. Despite these considerations, the empirical evidence obtained in this study did not reveal any active instances of the primary vulnerability associated with shared RSA prime factors. Table 1 summarizes the key differences between the findings of the original study and those of the present investigation.

TABLE 1. COMPARISON WITH THE ORIGINAL STUDY.

Aspect	Heninger et al. [1]	This Study (2025)
Dataset size	Millions of keys	1,444 certificates
Scope	Global	Latin American domains
Vulnerability observed	Widespread	None detected
Factorable RSA keys	≈0.5%	0%

Repeated RSA keys	60–65%	3.32%

DISCUSSION

The absence of RSA moduli sharing prime factors in this dataset contrasts sharply with the findings reported in 2012. Several factors may explain this difference.

First, modern operating systems have significantly improved entropy collection mechanisms, particularly during system boot. Second, contemporary hardware frequently incorporates dedicated random number generators, reducing reliance on potentially weak software-based entropy sources. Third, certificate authorities have adopted stricter cryptographic policies, including minimum key sizes and rejection of insecure parameters.

The relatively low percentage of repeated RSA moduli observed (3.32%) further suggests improved key management practices compared to earlier large-scale scans, where default or duplicated keys were common in embedded devices.

Nevertheless, the limited dataset size restricts the generalizability of these findings. While no vulnerabilities were observed in this sample, isolated weaknesses may still exist in legacy systems, unmaintained devices, or constrained embedded environments.

CONCLUSIONS

This study conducted an empirical evaluation of RSA common-factor vulnerabilities in contemporary TLS certificates associated with Latin American domains. By replicating the GCD-based methodology introduced by Heninger et al., no evidence of shared prime factors among RSA moduli was found.

The results suggest that the specific entropy-related RSA vulnerabilities identified more than a decade ago have been largely mitigated in modern certificate infrastructures. Improvements in randomness generation, cryptographic libraries, and certificate authority practices appear to have significantly strengthened RSA key generation.

Despite these positive findings, continuous cryptographic auditing remains essential. Expanding the dataset, including broader geographic scopes and legacy systems, would provide additional assurance that RSA-based infrastructures remain secure in practice.

AUTHOR CONTRIBUTION

The author's contributions to this article are as follows:

Leonardo Lizcano Pinto: Development, Methodology, Software, Validation, Formal analysis, Data curation, and Visualization.

Daniela Ospino-Balcázar: Writing–original draft, Writing–review and editing.

The author reviewed the results and approved the final version of the article.

CONFLICT OF INTERESTS

The authors declare that they have no interests or financial relationships that could have influenced this work.

REFERENCES

[1] N. Heninger, Z. Durumeric, E. Wustrow, and J. A. Halderman, "Mining your {Ps} and {Qs}: Detection of Widespread Weak Keys in Network Devices," in Proceedings of the 21st USENIX Security Symposium, USENIX Association, 2012, pp. 35–52.

[2] S. Yilek, E. Rescorla, H. Shacham, B. Enright, and S. Savage, "When Private Keys are Public: Results from the 2008 {Debian} {OpenSSL} Vulnerability," in Proceedings of the ACM Internet Measurement Conference (IMC), 2009, pp.

15–27. doi: 10.1145/1644893.1644896.

[3] M. Hastings, J. Fried, and N. Heninger, “Weak Keys Remain Widespread in Network Devices,” in Proceedings of the 2016 Internet Measurement Conference, 2016. doi: 10.1145/2987443.2987486.

[4] J. Kilgallin and R. Vasko, “Factoring {RSA} Keys in the {IoT} Era,” in 2019 First IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications (TPS-ISA), 2019, pp. 184–189. doi: 10.1109/tps-isa48467.2019.00030.

[5] N. Samarasinghe and M. Mannan, “Another Look at {TLS} Ecosystems in Networked Devices vs. Web Servers,” Comput. Secur., vol. 80, pp. 1–13, 2019, doi: 10.1016/j.cose.2018.09.001.

[6] F. Nezhadian, E. Branca, A. Barzolevskaia, A. Natadze, and N. Stakhanova, “Measuring and Characterizing Propagation of Reuse {RSA} Certificates and Keys Across {PKI} Ecosystem,” IEEE Transactions on Networking, vol. 33, pp. 595–611, 2025, doi: 10.1109/tnet.2024.3495617.

Leonardo Lizcano Pinto. Master's student in systems and computer engineering at the Universidad del Norte, Barranquilla, Colombia.
Daniela Ospino-Balcázar. Master's student in systems and computer engineering at the Universidad del Norte, Barranquilla, Colombia.