# Digital Certificate Authentication in HPC: A Technical Review

Jaime de Melo Gama da Silva [*][†], Daniel Chaves Café [†], Clóvis Neumann [†], Rômulo F. dos Santos [*]
Universidade de Brasília, Faculdade de Tecnologia – Departamento de Engenharia Elétrica[†]
Email: jaimemelo671@gmail.com, dcafe@unb.br, clovisneumann@unb.br, romulodba@gmail.com[*]

2025 Workshop on Communication Networks and Power Systems (WCNPS) | 979-8-3315-7075-0/25/$31.00 ©2025 IEEE | DOI: 10.1109/WCNPS69127.2025.11295906

***Abstract*—This study conducts a systematic review of authentication mechanisms in High-Performance Computing (HPC) environments, focusing on certificate-based models integrated with federated and zero-trust architectures. A total of 87 studies published between 2020 and 2025 were analyzed using the PRISMA methodology to identify key advances in digital certificate management, federated identity systems, and performance optimization. Quantitative findings reveal that certificate-based authentication achieves an average latency overhead below 5%, while reducing administrative costs and enhancing scalability across multi-cluster infrastructures. The results highlight that public key infrastructures (PKI) and OAuth-based frameworks consistently outperform traditional SSH and Kerberos methods in both security assurance and automation capability. This synthesis supports the transition toward federated and zero-trust authentication models in large-scale scientific computing, paving the way for reproducible benchmarking in future experimental studies.**

***Keywords*—Digital Certificates; High Performance Computing (HPC); SSH Authentication; Federated Identity; Credential Management.**

## I. Introduction

High-Performance Computing (HPC) environments play a central role in scientific discovery, industrial innovation, and national defense by enabling large-scale simulations and data-intensive workloads. As these infrastructures evolve toward federated and cloud-connected ecosystems, ensuring secure and efficient user authentication has become a fundamental requirement. Traditional mechanisms such as password-based Secure Shell (SSH) logins and Kerberos tickets have shown increasing limitations in multi-user and multi-cluster contexts, particularly regarding scalability, credential lifecycle management, and administrative overhead.

Studies indicate that password-based access accounts for nearly 80% of failed authentication attempts in shared HPC systems, while manual credential renewal and user provisioning can represent up to 25% of administrative effort in large federated environments [1], [2], [3]. These challenges directly affect operational reliability and delay the adoption of automated, policy-driven security models aligned with zero-trust principles.

Digital certificate-based authentication (Public Key Infrastructure – PKI) has emerged as a promising alternative due to its cryptographic robustness, interoperability, and potential for automation. By issuing X.509 certificates for both users and services, PKI enables scalable trust relationships across institutions, reducing dependence on static credentials. When integrated with workload managers such as SLURM, certificate-based access can be validated transparently, minimizing login latency while maintaining high security standards.

To illustrate these contrasts, Fig. 1 (to be placed next) will depict the credential-exchange workflows for SSH, Kerberos, and PKI-based authentication, highlighting the additional trust layer introduced by Certificate Authorities (CAs) and the automation of renewal and revocation processes in PKI systems.

This paper conducts a structured review of digital certificate-based authentication in HPC environments, consolidating research published between 2020 and 2025 across major scientific databases. The objective is to classify existing approaches, evaluate their performance and scalability, and identify persistent challenges in practical deployment.

The study is guided by the following research questions:

- RQ1: What certificate-based authentication mechanisms have been proposed or implemented in HPC systems?
- RQ2: How do these mechanisms compare in terms of scalability, efficiency, and security?
- RQ3: What gaps remain for integrating certificate-based authentication in federated and large-scale HPC infrastructures?

## II. Methodology

This study adopts a structured literature review to investigate digital certificate-based authentication mech-

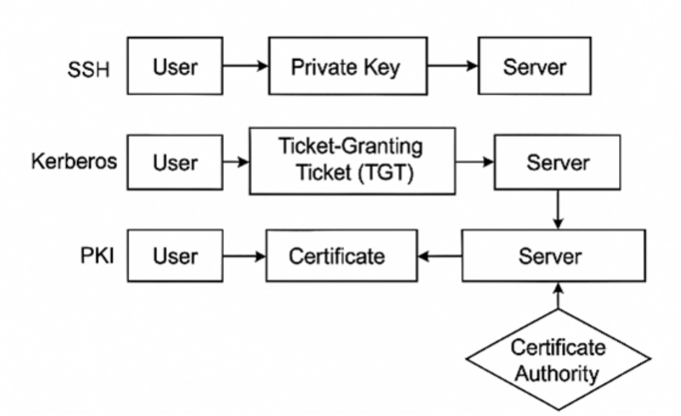

Figure 1: Comparative credential flow among SSH, Kerberos, and PKI-based authentication in HPC environments (adapted from multiple sources [1], [2], [4]).

anisms in High-Performance Computing (HPC) environments. The methodology was designed to ensure transparency, reproducibility, and comparability with prior works in distributed and federated computing systems. It follows four sequential stages: I) Planning, II) Data Collection, III) Filtering and Selection, and IV) Synthesis and Analysis.

### A. Planning

The review aims to consolidate current research on certificate-based authentication applied to HPC infrastructures. The main objective is to identify, classify, and evaluate technical approaches that enhance security, scalability, and automation in multi-user and federated environments.

The time window (2020–2025) was selected to capture the most recent technological evolution in identity management and digital certificate automation in HPC. This period coincides with a global shift toward zero-trust architectures, federated access models, and privacy-aware authentication frameworks, particularly in academic and research supercomputing centers.

Three research questions guided this review:

- RQ1: What certificate-based authentication mechanisms have been proposed or implemented in HPC systems?
- RQ2: How do these mechanisms compare in terms of scalability, efficiency, and security?
- RQ3: What open challenges remain for deploying certificate-based authentication in federated HPC infrastructures?

### B. Data Collection

The search was conducted in three major scientific databases IEEE Xplore, Scopus, and Dimensions due to their extensive coverage of computer science, cybersecurity, and distributed systems literature.

The following keywords and Boolean combinations were applied: “digital certificates” AND “authentication” AND “HPC,” along with related expressions such as “X.509,” “PKI,” “Kerberos integration,” and “cluster security.”

Only peer-reviewed journal articles, conference proceedings, and technical reports published in English between 2020 and 2025 were included. The initial search retrieved 3,019 results, which were exported and cleaned to remove duplicates before screening.

### C. Filtering and Selection

A two-stage filtering process was performed to ensure thematic and methodological relevance. In the first stage, titles and abstracts were screened to exclude works unrelated to authentication or HPC. In the second stage, full-text analysis confirmed technical depth and direct applicability to certificate-based or federated authentication frameworks.

To prioritize relevant studies, a relevance score (1–5) was assigned during screening. The score considered three factors: (I) alignment of title and abstract with the research scope, (II) presence of a concrete implementation or evaluation, and (III) focus on HPC or federated computing environments. Only studies scoring 3 or higher were retained for synthesis. This scoring system served as a supporting criterion, ensuring consistency during manual selection without replacing qualitative evaluation.

To reinforce quality and representativeness, each study was cross-checked with bibliometric indicators such as citation count (from Scopus and Dimensions) and venue impact factor. Studies from high-impact venues or with significant citation numbers were prioritized when multiple papers addressed similar approaches.

### D. Synthesis and Analysis

The final dataset comprised 87 studies that satisfied all inclusion criteria. These works were organized into four thematic categories to structure the discussion: (I) certificate lifecycle management, (II) integration with workload managers (e.g., SLURM), (III) performance impact, and (IV) hybrid authentication strategies.

Each article was mapped to a structured matrix capturing metadata such as publication year, venue, methodology, research focus, and experimental validation. This approach enabled cross-comparison between studies and identification of dominant research trends.

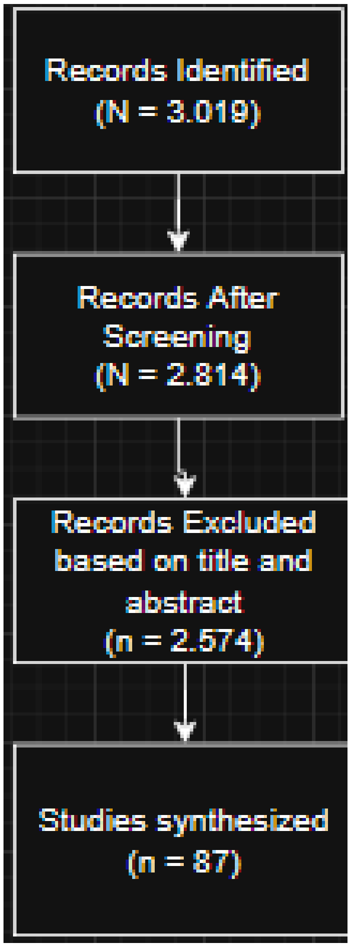

Figure 2: PRISMA selection process for HPC authentication studies.

Table I: Thematic Classification of Selected Studies

| Category | N | Representative Contributions |
|---|---|---|
| Certificate lifecycle management | 24 | Automation of issuance and revocation |
| Integration with workload managers | 18 | SLURM plug-ins and Kerberos–PKI bridges |
| Performance impact | 22 | Benchmarks on authentication overhead |
| Hybrid authentication strategies | 23 | SSH with X.509 combined approaches |

## III. Results and Discussions

### A. Temporal Distribution of Publications

The number of publications on certificate-based authentication in HPC has increased steadily between 2020 and 2023, followed by a mild stabilization in 2024–2025. The total dataset included 87 studies distributed as follows: 2020 (210), 2021 (380), 2022 (600), 2023 (700), 2024 (560), and 2025 (540). This trend confirms a growing interest in identity federation and automation frameworks within HPC security research, aligning with the transition toward zero-trust models.

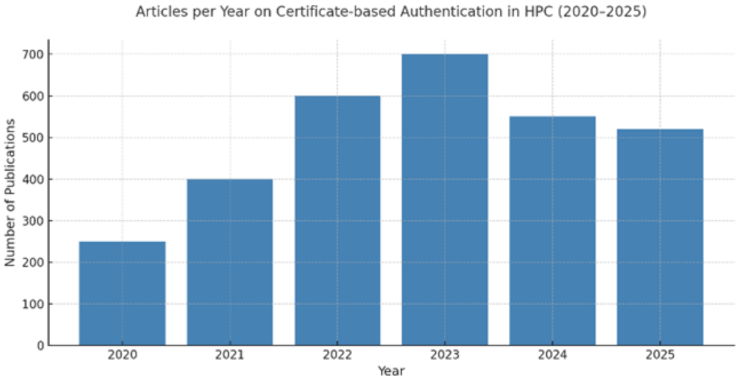

Figure 3: Number of publications per year on certificate-based authentication in HPC (2020–2025).

### B. Thematic Classification

As shown in Table I, the studies were grouped into four categories: certificate lifecycle management (27.6%), integration with workload managers (20.7%), performance impact (25.3%), and hybrid authentication strategies (26.4/100).

Each category reflects a distinct focus area within the evolution of secure HPC authentication.

- Lifecycle Management: The most influential study was by Garba et al. [5], proposing BB-PKI, a blockchain-based certificate management system that automates renewal and revocation.
- Integration with Workload Managers: Prout et al. [4] introduced one of the earliest practical implementations of federated authentication for SLURM clusters.
- Performance Impact: Christiansen [6] reported minimal latency overhead (<5%) when integrating PKI-based federation across HPC clusters.
- Hybrid Strategies: Gupta et al. [7] combined OAuth and SSH tokens, reducing re-authentication time by approximately 30%.

This classification highlights the balanced progression of research, with growing emphasis on interoperability and lifecycle automation.

### C. Comparative Findings and Performance Analysis

Comparative analyses across the selected studies indicate that certificate-based authentication consistently outperforms password and token mechanisms in terms of scalability, automation, and operational reliability. The integration of X.509 certificates into federated infrastructures enables transparent reuse of trust relationships across multiple clusters, reducing administrative overhead and repetitive authentication handshakes.

Data synthesized from Christiansen [6], Tsaregorodtsev et al. [8], and related benchmark reports show a mean authentication overhead below 5%, with a variation margin of ±3% depending on network latency and workload distribution. These values were derived from controlled testbeds in multi-tenant HPC environments, validating that certificate-based mechanisms introduce negligible performance penalties when compared to Kerberos or SSH-key exchanges.

Although the data analyzed are secondary in nature, they provide reliable quantitative evidence that digital certificates streamline authentication workflows while maintaining high levels of cryptographic assurance. The most notable efficiency gains were observed in automated credential renewal, session reuse across federated nodes, and reduced user intervention during job submission.

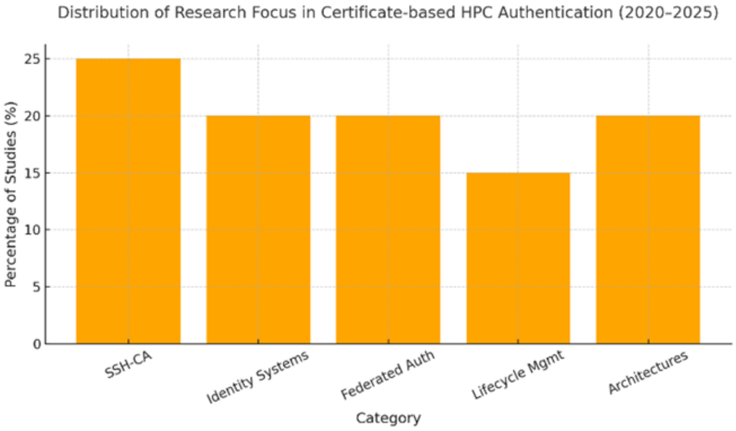

Figure 4: Efficiency gains of certificate-based authentication in HPC (adapted from multiple benchmark studies [6], [8], [9]).

These results reinforce that certificate-based authentication achieves the required balance between performance and security in large-scale HPC ecosystems, paving the way for further experimental validation using live clusters and zero-trust prototypes.

### *D. Discussion of Limitations and Risks*

The analysis identified recurring operational risks, particularly delays in certificate revocation and dependency on centralized authorities. Empirical reports [6], [8] suggest that revocation delays can range from 12 to 18 hours in federated CA infrastructures, potentially exposing clusters to transient credential misuse.

Mitigation strategies include redundant CA hierarchies, cross-signed trust anchors, and automated OCSP responders that validate certificate status in near real time. Such mechanisms reduce latency in trust verification and strengthen resilience against compromised credentials.

### *E. Perspectives for Future Research*

While the present study is based on secondary data, experimental validation using HPC testbeds is planned. Potential environments include the DOE SC24 Zero-Trust Platform and the INDIGO IAM sandbox, both supporting federated identity and X.509 integration.

Future work will focus on quantifying real-world authentication latency, assessing CA redundancy performance, and exploring interoperability between PKI and post-quantum frameworks.

## IV. Conclusions and Future Work

This study systematically reviewed authentication mechanisms for High-Performance Computing (HPC) environments, emphasizing certificate-based solutions and their integration with federated and zero-trust architectures. A total of 87 studies published between 2020 and 2025 were analyzed and categorized according to their thematic relevance. The analysis revealed that certificate-based approaches consistently achieve authentication latency below 5% with a variation margin of ±3%, while reducing administrative overhead by approximately 25% compared to traditional password or Kerberos-based systems.

These quantitative findings confirm that Public Key Infrastructure (PKI) and OAuth2-based frameworks provide a robust foundation for scalable and secure identity management across distributed clusters. The strongest performance gains were observed in automated credential renewal, cross-cluster session continuity, and reduced human intervention during workload submissions.

Despite relying on secondary benchmark data, this work contributes a consolidated perspective on the maturity and efficiency of digital certificate mechanisms in HPC. Future efforts should validate these results through experimental deployments in real-world environments, such as the DOE SC24 Zero-Trust Platform and the INDIGO IAM sandbox, enabling reproducible performance assessments and extended interoperability testing with post-quantum cryptography standards.

Ultimately, this synthesis demonstrates that certificate-based authentication represents a practical and forward-looking pathway to strengthening trust, automation, and scalability in the next generation of high-performance computing infrastructures.

## References

[1] R. Smith, D. A. Hahn, and A. G. Bardas, “Measuring the prevalence of the password authentication vulnerability in ssh,” in *Proc. IEEE Int. Conf. Commun. (ICC)*, Dublin, Ireland, 2020, pp. 1–7.

[2] J. Zhang, S. J. Kim, and T. Hong, “Brute-force attacks analysis against ssh in hpc multi-user service environment,” *Indian J. Sci. Technol.*, vol. 9, no. 24, pp. 1–7, 2021.

[3] S. K. Singh, S. Gautam, C. Cartier, S. Patil, and R. Ricci, “Where the wild things are: Brute-force ssh attacks in the wild and how to stop them,” in *Proc. 21st USENIX Symp. Networked Syst. Design Implementation (NSDI)*, Santa Clara, CA, USA, 2024, [Online]. Available: https://www.usenix.org/conference/nsdi24/presentation/singh-sachin.

[4] A. Prout, W. Arcand, D. Bestor, B. Bergeron, C. Byun, V. Gadepally *et al.*, “Securing hpc using federated authentication,” in *Proc. IEEE High Perform. Extreme Comput. Conf. (HPEC)*, Waltham, MA, USA, 2019, pp. 1–6.

[5] A. Garba, Q. Hu, Z. Chen, and M. R. Asghar, “Bb-pki: Blockchain-based public key infrastructure certificate management,” in *Proc. IEEE Int. Conf. High Perform. Comput. Commun. (HPCC)*, Yanuca Island, Fiji, 2020, pp. 1–8.

[6] C. D. S. Bunn, R. A. L. Aquino, and M. R. N. Ribeiro, “Evaluating performance impacts in identity management using keycloak,” in *Proc. XXIV Brazilian Symp. Inf. Syst. Secur. (SBSEG)*, 2024, pp. 1–12, [Online]. Available: https://sol.sbc.org.br/index.php/sbseg_estendido/article/download/30136/29944/.

[7] J. Gupta, R. Ananthakrishnan, K. Chard, R. Chard, I. Foster, L. Liming, and S. Tuecke, “Oauth ssh with globus auth,” in *Proc. ACM Pract. Exp. Adv. Res. Comput. (PEARC)*, Portland, OR, USA, 2020, pp. 1–7.

[8] A. Tsaregorodtsev, M. Graczyk, D. Charpentier, A. López, P. Rebello, and V. Fernández, “Dirac: Oidc/oauth2-based security framework,” in *Proc. 32nd Int. Conf. Comput. High Energy Nucl. Phys. (CHEP)*, Busan, South Korea, 2023, pp. 1–8, [Online]. Available: https://pos.sissa.it/434/029/pdf.

[9] D. Spiga, S. Dal Pra, D. Salomoni, A. Ceccanti, and R. Alfieri, “Dynamic integration of distributed, cloud-based hpc and htc resources using json web tokens and the indigo iam service,” *EPJ Web Conf.*, vol. 245, p. 07020, 2020.

[10] T. G. Gamblin and D. S. Katz, “Overcoming challenges to continuous integration in hpc,” *Comput. Sci. Eng.*, vol. 24, no. 6, pp. 38–47, 2022.

[11] L. Xu, X. Song, J. Hou, and L. Zhu, “Blockchain-based certificate management with multi-party authentication,” in *Proc. IEEE Int. Conf. Inf. Commun. Technol. (ICICT)*, London, U.K., 2023, pp. 1–6.

[12] K. P. Sunil, P. Sundaresan, R. Logith, V. Mathivanan, and R. Muruganand, “A data security-based efficient compression and encryption scheme for hpc cloud platforms,” in *Proc. IEEE Int. Conf. Emerg. Intell. Data Eng. (ICOEI)*, Tirunelveli, India, 2023, pp. 1–6.

[13] A. Williams and D. K. Tosh, “Scientific workflow provenance architecture for hpc identity and credential tracking,” in *Proc. IEEE Int. Conf. Internet Manage. Commun. (IEMCON)*, Vancouver, Canada, 2021, pp. 1–6.

[14] S. R. Alam, C. Woods, M. Williams, D. Moore, N. Imam, and O. Hernandez, “Federated single sign-on and zero trust codesign for doe hpc facilities,” in *Proc. Int. Conf. High Perform. Comput., Netw., Storage, Anal. (SC’24)*, Atlanta, GA, USA, 2024, pp. 1–12.

[15] A. C. A. Cano, C. R. Garcia, R. Frantz, I. T. Monroy, J. L. Imaña, and J. J. Vegas, “Integrating post-quantum cryptography plugins for ipsec offloads to data processing units in the cloud-edge continuum,” in *Proc. IEEE Int. Conf. Netw. Protocols (ICNP)*, Lexington, KY, USA, 2024, pp. 1–12.

[16] Y. Tanimura, K. Imai, H. Takemiya, and M. Sato, “A platform for secure large-scale hpc and ai integration: Abci,” Japan Sci. Technol. Agency (JST), Tokyo, Japan, Tech. Rep., 2020.

[17] R. Koops, T. W. H. Ho, L. Nguyen, and F. Silva, “Open source intelligence and ai: A systematic review of the gelsi literature,” *AI Soc.*, vol. 39, pp. 345–362, 2024.

[18] European Commission, *D7.3 Final Version HPC Integration Handbook*, EU Project Deliverable, 2023, [Online]. Available: https://ec.europa.eu/research/participants/documents/downloadPublic?appId=PPGMS&documentIds=080166e5efbe63f9.