Twenty-Five Years of the Brazilian Cybersecurity Symposium (SBSeg): A Retrospective Review

Silvio E. Quincozes[1,2], Camilla B. Quincozes[1,2], Paulo Souza[1]

[1]AI Horizon Labs, Universidade Federal do Pampa (UNIPAMPA)

[2]PPGCO – Universidade Federal de Uberlândia (UFU)

silvioquincozes@unipampa.edu.br, sequincozes@ufu.br

camillaquincozes@ufu.br, paulosila@unipampa.edu.br

Abstract. *The Brazilian Cybersecurity Symposium (SBSeg) reached its twenty-fifth edition in 2025, reflecting the evolution of cybersecurity research in Brazil. This paper presents an article-level longitudinal survey of the SBSeg main track from 2001 to 2025, covering 25 proceedings volumes and 824 papers. We analyze the corpus through three axes: community formation, research agenda, and methodological maturity. Specifically, we examine who published in the venue, which topics structured its scientific agenda, and how research practices evolved across prototypes, datasets, empirical evaluation, formal methods, machine learning, vulnerability analysis, and user-centered studies. The results show SBSeg's transition from information and computer systems security to a broader cybersecurity forum, preserving foundational areas while incorporating new waves such as privacy, IoT, blockchain, post-quantum cryptography, adversarial machine learning, and generative AI.*

1. Introduction

Scientific venues are more than publication outlets: they are longitudinal records of how research communities define problems, adopt methods, and respond to technological and societal change. In cybersecurity, this historical perspective is particularly relevant because the field evolves through successive waves of attacks, infrastructures, regulations, platforms, and defensive paradigms. A venue's publication record can therefore reveal not only how much a community has grown, but also which topics persisted, which research fronts emerged, which methods became dominant, and how the community reorganized itself around new cybersecurity challenges.

The Brazilian Cybersecurity Symposium (SBSeg), formerly known as the Brazilian Symposium on Information and Computer Systems Security, is the main archival venue for cybersecurity research in Brazil. Since its origins as the Workshop on Computer Systems Security, SBSeg has accompanied the maturation of the Brazilian cybersecurity community across a quarter century. Its proceedings capture the transition from classical information and systems security concerns, such as cryptography, authentication, access control, network security, and digital forensics, toward a broader cybersecurity agenda that now includes privacy, IoT (Internet of Things), cyber-physical systems, blockchain, post-quantum cryptography, adversarial machine learning, and generative AI(Artificial Intelligence).

Despite this trajectory, SBSeg has not yet been examined as an article-level longitudinal corpus. Existing views of the event are typically implicit, fragmented, or restricted to specific editions, topics, or communities. Counting papers per year is therefore insufficient to explain how the venue evolved. A stronger analysis must connect three dimensions: who published at SBSeg, what the community studied, and how its methodological standards evolved.

This paper presents an article-level longitudinal survey of SBSeg research from 2001 to 2025 through three complementary lenses. *Community formation* examines authors, institutions, recurrence patterns, and institutional participation over time. *Research agenda* analyzes titles, keywords, representative papers, and thematic categories to identify persistent areas and emerging fronts. *Methodological maturity* investigates how SBSeg papers produced evidence, including prototypes, datasets, empirical evaluations, formal models, simulations, AI models, vulnerability analyses, and user-centered studies.

This study is guided by the following research questions:

- **RQ1. Research agenda:** Which cybersecurity topics persisted, expanded, declined, or emerged across the 25-year trajectory of SBSeg?
- **RQ2. Community formation:** How did the authorship, institutional participation, and collaboration structure of SBSeg evolve from 2001 to 2025?
- **RQ3. Historical interpretation:** What does the co-evolution of topics and community reveal about the maturation and future of cybersecurity research in Brazil?

Our contribution is fourfold. First, we construct an article-level longitudinal dataset of SBSeg main-track papers from 2001 to 2025, consolidating metadata on years, authors, institutions, titles, keywords, and derived analytical labels. Second, we propose and apply a multi-axis taxonomy that organizes SBSeg research by community, thematic area, methodological approach, and security concern. Third, we provide longitudinal analyses of institutional participation, author productivity, topic evolution, first appearances of research fronts, and methodological trends. Fourth, we identify representative SBSeg papers and use them to build a critical narrative of how Brazilian cybersecurity research moved from foundational information and computer systems security toward a broader and more methodologically diverse cybersecurity agenda.

2. Methodology

The corpus comprises all 824 main-track SBSeg papers published from 2001 to 2025. We excluded minicourses, tutorials, workshops, panels, prefaces, editorials, demos, and satellite events, while including both full and short papers published in the main proceedings. For each paper, we collected the edition year, paper ID, title, authors, institutional affiliations, and author-provided keywords when available.

The dataset was built through a semi-automated metadata extraction pipeline followed by manual cleaning, normalization, and entity resolution. Author-name variants caused by abbreviations, surname conventions, diacritics, and inconsistent citation formats were consolidated when the available metadata supported the match. Institutional acronyms, full names, and affiliation variants were similarly harmonized. To avoid inflated counts, each normalized author and institution was counted at most once per paper.

The thematic analysis combined author-provided keywords with standardized descriptors derived from paper titles, since keywords were absent or inconsistent in several

editions. Recurring title expressions were grouped into broader topics, and papers could be assigned to more than one topic when their titles supported multiple themes. Generic standalone terms, such as "security" and "attack", were excluded from the thematic ranking.

ChatGPT 5.5 Plus supported terminology harmonization, candidate entity matching, title-pattern identification, and consistency checking. It was used as an assistance mechanism rather than as an autonomous classifier: all suggested author or institution matches and all title-derived descriptors were manually reviewed, corrected when necessary, and validated by the authors against the original metadata before being included in the analyses. Although this process improves consistency across heterogeneous proceedings, subjective decisions may remain in entity resolution and thematic grouping.

3. Research Themes Over Time: What does SBSeg study?

To characterize the evolution of SBSeg's research agenda from 2001 to 2025, we derived thematic keywords directly from article titles and grouped them into broader research categories. This title-based procedure complements author-provided keywords by capturing topics even when metadata is absent or inconsistent across editions. Overall, 824 articles were processed, 576 matched at least one title-derived grouped keyword, and 190 matched more than one grouped keyword. Each article could contribute to multiple groups when supported by its title, but at most once per group in a given year. The top 15 groups were then selected by article volume, excluding overly generic standalone terms such as "security" and "attack".

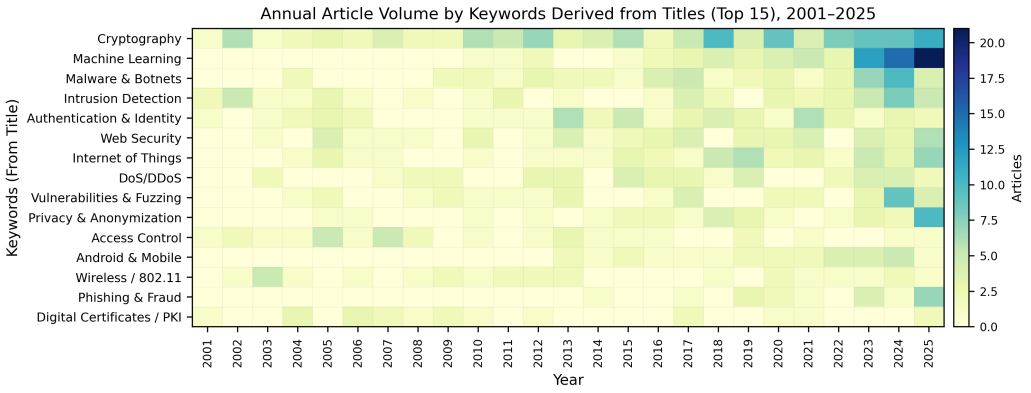

Figure 1. Top 15 keywords derived from article metadata (grouped by similarity)

Figure 1 shows that SBSeg's research agenda combines persistent foundational topics with newer research fronts. Cryptography, authentication, access control, and intrusion detection appear across much of the series, reflecting long-standing concerns with secure communication, identity, authorization, and defensive monitoring. From the 2010s onward, the agenda becomes more diversified, with growing visibility of machine learning, malware and botnet analysis, Internet of Things, web security, privacy, vulnerabilities, phishing, Android/mobile security, and DoS/DDoS. This suggests a cumulative expansion of SBSeg from a venue centered on classical computer and network security toward a broader cybersecurity forum shaped by emerging platforms, data-driven methods, and application-layer threats. In the following subsections, we detail the SBSeg research history.

3.1. Workshop on Computer Systems Security (2001–2004)

The first four editions, still organized as the *Workshop on Computer Systems Security*, reveal a formative stage in which SBSeg was primarily anchored in computer systems, networks, cryptography, intrusion detection, and access control. The 2001 and 2002 editions already established several threads that would remain recurrent in the following years: cryptographic protocols, defensive monitoring, authentication, authorization, and secure distributed administration [Campello et al. 2001, Ribeiro and Custódio 2002, Motta and Furuie 2002]. They also introduced applied concerns that would later become persistent in the venue, such as electronic voting, digital forensics, and security for distributed or institutional systems. Overall, this early agenda suggests a movement from isolated protection mechanisms toward security as an operational infrastructure, in which protocols, monitoring, identity, and access-control mechanisms were increasingly treated as components of larger managed environments.

The 2003 edition marks the first clear thematic wave of the workshop: wireless and mobile network security. Topics such as IEEE 802.11 wireless local area networks, WEP , VPN/IPSec, OpenBSD-based access points, MPLS , virtual private networks, and ad hoc networks dominated the program, showing that the community was responding to the practical risks introduced by emerging wireless and virtualized communication infrastructures [Peres and Weber 2003]. At the same time, recurring themes such as cryptography, access control, authentication, electronic voting, and IDS evaluation remained active, suggesting continuity beneath the yearly thematic shift. In 2004, the agenda broadened again toward formalized security services, operational evidence, policy-aware management, and secure distributed infrastructures, including SPKI/SDSI, X.509 digital certificates, OTPs , RBAC , remote VPN access, ad hoc routing, WSNs , IP traceback, honeypots, Sec-SLA, and virtual-machine-based IDS protection [Laureano et al. 2004].

Taken together, these four editions portray the workshop as a formative space in which Brazilian security research began to move from mechanism-centered studies to broader views of security as an infrastructural and socio-technical concern. The recurring presence of cryptography, authentication, access control, intrusion detection, and network security provided the workshop's initial backbone, while the early appearance of topics such as voting, forensics, wireless networks, sensor networks, honeypots, and security management anticipated research fronts that would gain relevance in later editions. Thus, the 2001–2004 period should not be read merely as a preliminary phase, but as the foundation of SBSeg's long-term identity: a venue rooted in computer systems security, yet already open to distributed infrastructures, operational evidence, and applied cybersecurity problems.

3.2. When Security Became Infrastructure

The 2005–2009 period marks the transition from the workshop's formative agenda to the early identity of SBSeg as a symposium. If the 2001–2004 phase was rooted in mechanisms for protecting systems, networks, and protocols, the following quinquennium shows security becoming an infrastructural concern. In 2005, this shift was already visible in the combination of intrusion detection, alert correlation, policy-based access control, Web Services, grids, peer-to-peer systems, and wireless sensor networks. Rather than treating security as an isolated component, the community increasingly

explored how mechanisms could be composed, managed, and embedded into distributed environments, as illustrated by works on large-scale security configuration management, IDS composition through Web Services, and secure wireless sensor networks [de Albuquerque et al. 2005, Brandão et al. 2005, Oliveira et al. 2005].

Across the quinquennium, some topics became recurrent backbones: cryptography, access control, intrusion detection, network security, and secure distributed systems. However, their role changed over time. Cryptography remained central, but appeared not only in primitive design or cryptanalysis, but also as a foundation for PKI (Public Key Infrastructure), digital signatures, secure voting, WSNs (Wireless Sensor Networks), and hardware-supported ceremonies. Access control evolved toward policy enforcement, UCON (Usage Control), RBAC (Role-Based Access Control), SoD (Separation of Duty), trust chains, and compliance, reflecting a growing concern with authorization as an enforceable and manageable service in distributed environments [Silva et al. 2007]. Intrusion detection also matured from standalone detectors to heterogeneous compositions, alert correlation, attack traceback, IDS (Intrusion Detection System) interoperability, and operational monitoring. This period also introduced or consolidated themes that would later gain importance, such as privacy for LBSs (Location-Based Services), secure software development, DRM (Digital Rights Management) interoperability, trusted platforms, security ontologies, honeypots, vulnerability management, and adaptive security services.

The most distinctive transformation appeared in 2008 and 2009, when SBSeg began to display a more measurement-oriented and evidence-driven agenda. Topics such as network-flow storage, passive DNS (Domain Name System) analysis, anomaly filtering, malware taxonomy, obfuscated binary analysis, embedded-software integrity verification, and user-centered information-security management indicate a methodological broadening beyond protocol design and formal modeling. This transition is exemplified by works on network-flow storage for traffic and security analysis and on malware taxonomy based on automatically propagated threats [Corrêa et al. 2008, Ceron et al. 2009]. Thus, the early symposium years should be read as the moment when SBSeg's agenda expanded from securing individual mechanisms to understanding security as a managed ecosystem: distributed, policy-aware, operational, and increasingly supported by empirical evidence.

3.3. The Rise of Cybersecurity Ecosystems (2010–2014)

The 2010–2014 quinquennium marks SBSeg's expansion from secure mechanisms and distributed infrastructures toward a broader cybersecurity ecosystem. Cryptography remained a recurring backbone, covering public-key schemes, oblivious transfer, hash functions, timestamps, lattice-based constructions, side-channel countermeasures, post-quantum signatures, PUFs (Physical Unclonable Functions), authenticated encryption, and secure protocols for constrained devices [Oliveira and López 2013]. Yet its role shifted: cryptographic research increasingly coexisted with applied problems in wireless networks, cloud computing, mobile platforms, smart grids, e-health, electronic voting, and identity management. Therefore, in this period, SBSeg preserved its formal roots while expanding into emerging infrastructures.

A second defining feature of the period was the rise of operational, empirical, and data-driven security. From 2010 onward, intrusion detection, anomaly detection, alert correlation, botnet mitigation, malware analysis, spam detection, traceback, and

incident handling became increasingly prominent. The agenda moved beyond proposing isolated detectors to exploring classification, time-series modeling, attack correlation, automated response, IDS labeling, malware behavior monitoring, web anomaly detection, and security assessment of real platforms. This shift is especially visible in works on machine-learning-based intrusion detection, Android malware analysis, and anomaly monitoring in cloud environments, which suggested a community increasingly concerned with measurable evidence, operational telemetry, and scalable defense mechanisms [Henke et al. 2011, Afonso et al. 2013, Ferreira and de Geus 2013].

The quinquennium also consolidated SBSeg as a venue attentive to the changing shape of digital infrastructure. Web security, Android malware, static vulnerability detection, software protection, memory-violation mitigation, secure logging, ROP defense, and scanner assessment signaled the growing importance of software and application security. At the same time, cloud computing, Future Internet, cognitive radio, content-centric networking, IoT, Web of Things, smart grids, vehicular networks, 6LoWPAN, and NFC indicated that the research agenda was moving toward heterogeneous, mobile, virtualized, and cyber-physical environments [Cervantes et al. 2014]. In this sense, 2010–2014 represents the point at which SBSeg's identity becomes recognizably modern: still grounded in cryptography and network security, but increasingly shaped by empirical detection, software security, cloud/IoT infrastructures, identity management, and socio-technical applications.

3.4. The Rise of Modern Cybersecurity: IoT, SDN, Malware, and Machine Learning (2015–2019)

The 2015–2019 period marks the consolidation of SBSeg as a venue aligned with modern cybersecurity concerns. Cryptography remained a strong and recurring foundation, appearing in works on homomorphic encryption, lattice-based constructions, identity-based cryptography, PAKE protocols, post-quantum candidates, hash-based signatures, white-box cryptography, secure computation, and privacy-preserving mechanisms. However, the period was no longer dominated by cryptography as an isolated research front. Instead, cryptographic techniques increasingly appeared in connection with IoT, smart devices, digital voting, blockchain, trusted execution environments, encrypted analytics, and privacy-preserving data processing, as illustrated by work on identity-based cryptography from the perspective of Internet of Things environments [Neto et al. 2015]. This shift suggests that SBSeg's cryptographic tradition was being absorbed into broader systems and domains.

The most visible transformation of the quinquennium was the rise of programmable, industrial, and IoT-oriented security. SDN, P4, network-function chains, DDoS mitigation, IoT access control, CoAP fuzzing, MQTT security, industrial control networks, SCADA intrusion detection, smart metering, cyber-physical attacks, and provider-backbone defenses became central topics. These works indicate a community increasingly concerned with infrastructures that are dynamic, programmable, heterogeneous, and exposed to large-scale attacks, including SDN-based defenses against denial-of-service attacks, supervised-learning-based intrusion detection for SCADA networks, and IoT datasets for host-based intrusion detection research [Piedrahita et al. 2015, Vasquez et al. 2017, Bezerra et al. 2018]. At the same time, authentication and privacy evolved toward continuous authentication, biometrics, device

fingerprinting, web fingerprinting, trajectory anonymization, smart-meter privacy, and privacy-aware mobile and web systems, showing that security was expanding from network protection to persistent identity, behavioral traces, and user/device ecosystems.

A third defining feature of this period was the acceleration of data-driven and adversarial security research. Malware analysis, ransomware, phishing, social bots, cryptomining, botnet detection, web-attack detection, adversarial images, JavaScript taint analysis, ROP detection, binary instrumentation evasion, and forensic search became increasingly prominent. Machine learning moved from a methodological option to a recurring security instrument, supporting intrusion detection, malware classification, social bot detection, phishing analysis, and anomaly detection, while also becoming itself a target of adversarial manipulation. Work on evasive malware analysis illustrates this maturation of the malware research agenda and its growing attention to anti-analysis and adversarial behavior [Botacin et al. 2017]. Taken together, 2015–2019 represents the moment when SBSeg's agenda becomes recognizably contemporary: rooted in cryptography and network security, but increasingly organized around IoT/SDN infrastructures, privacy-preserving systems, malware ecosystems, trusted hardware, blockchain, and machine-learning-supported defense.

3.5. AI-Native Cybersecurity (2020–2025)

The 2020–2025 period marks the point at which SBSeg becomes clearly shaped by data-intensive and AI-supported cybersecurity. Machine learning, which had appeared in earlier editions as one methodological option among others, became a recurring organizing force for intrusion detection, malware analysis, phishing detection, botnet identification, fraud detection, vulnerability analysis, anomaly detection, and cyber-physical monitoring. This movement intensified across the period: 2020 emphasized ML-based detection, federated learning, Android traces, phishing, botnets, and malware; 2021 expanded toward adversarial learning, zero-shot detection, and deep models; 2023 consolidated explainable and transfer-learning approaches for malware and IDSs; and 2024 made XAI, adversarial robustness, poisoning defenses, synthetic data, and interpretable IDSs central to the agenda [de Camargo et al. 2021, Bragança et al. 2023, Mazetto and Zarpelão 2024]. By 2025, this trajectory had reached an AI-native stage, with works on synthetic-data-supported APT detection, concept-drift detection, Transformer-based host anomaly detection, IoT threat detection with Kolmogorov-Arnold Networks, smart-grid IDS feature enrichment, adversarial attacks against healthcare IDSs, Open RAN IDS evasion, and zero-day detection in 6G networks.

A second defining characteristic was the consolidation of identity, privacy, decentralization, and post-quantum readiness as strategic concerns. Privacy-preserving computing, anonymization, LGPD-oriented techniques, homomorphic computing, differential privacy, browser fingerprinting, model inversion attacks, DP-SGD (Data Security with Differential Privacy), differential privacy synthetic data, semantic deduplication, and privacy-focused data commercialization emerged alongside decentralized identity, self-sufficient identity, Zero Trust, blockchain-based reputation, digital credentials, smart contracts, payment channel networks, and proof-of-stake consensus. Cryptography also remained active, but its emphasis increasingly shifted toward post-quantum schemes, hybrid cryptography, secure voting, Ring-LWE, Kyber/Saber, ML-KEM, FrodoKEM, NTRU, post-quantum AMI, post-quantum LoRaWAN security, hash-based signatures, blind sig-

natures, fair exchange, and post-quantum requirements for trusted platforms. Thus, the period reflects a landscape that responds not only to attacks but also to regulatory pressure, decentralized trust models, privacy-preserving analytics, and the anticipated transition to post-quantum infrastructures.

The 2025 edition adds a new layer to this trajectory: generative AI and LLM security. The most recent SBSeg agenda visibly converged around LLM-based incident categorization, RAG for forensic log analysis, prompt-based secure programming, prompt engineering for fake-news detection, LLM-assisted ModSecurity rule generation, NeMo Guardrails as an interaction firewall, modular guardrails for GenAI systems, universal jailbreaks, and data-extraction attacks [Paim et al. 2025]. At the same time, critical and emerging infrastructures remained central, with works on Open RAN, 6G, Kubernetes access control, UEFI firmware testing, RF cyber ranges, RPKI hardening, cloud monitoring, smart grids, WSNs, BRSKI, and critical IoT environments. Taken together, 2020–2025 represents the current frontier of SBSeg's research agenda: a mature cybersecurity forum organized around AI-supported defense, LLM and GenAI security, privacy-preserving learning, post-quantum transition, decentralized trust, and the protection of next-generation communication, IoT, and critical infrastructures.

3.6. Summary and Milestones

Figure 2 summarizes the thematic evolution of SBSeg from 2001 to 2025 by organizing the symposium's trajectory into successive research milestones. The early years were marked by foundational topics in systems and network security, including firewalls, VPNs, cryptography, intrusion detection, access control, e-voting, and digital forensics. Wireless security emerged as a prominent wave around 2003, reflecting the growing relevance of IEEE 802.11, WEP, VPN/IPSec, and ad hoc networks. Between 2005 and 2009, the agenda shifted toward security as infrastructure, with increasing attention to policy composition, operational evidence, honeypots, ontologies, and managed distributed architectures. From 2010 to 2014, SBSeg expanded toward a broader cybersecurity forum, incorporating malware, botnets, cloud computing, mobile security, smart grids, IoT, SDN, and empirical detection.

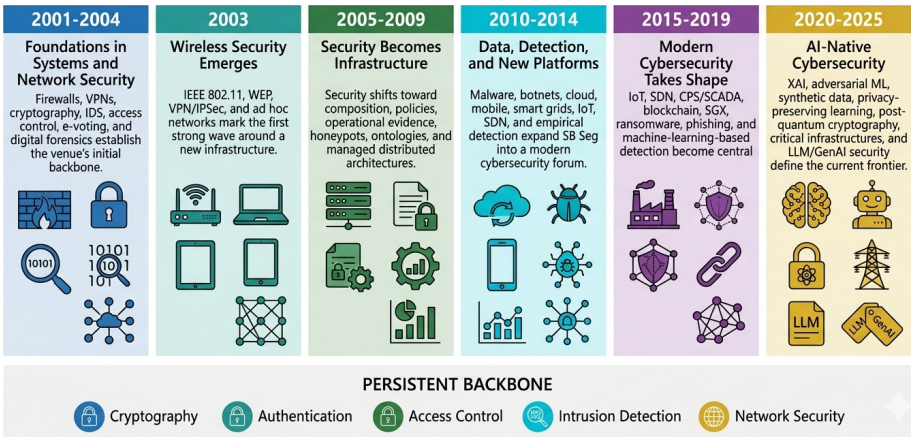

Figure 2. Thematic evolution of SBSeg from 2001 to 2025.

The most recent periods show the consolidation of modern and AI-native cybersecurity as central research fronts. From 2015 to 2019, topics such as IoT, SDN,

CPS/SCADA, blockchain, SGX, ransomware, phishing, and machine-learning-based detection became increasingly visible, indicating the symposium's alignment with contemporary cybersecurity challenges. From 2020 to 2025, the agenda moved further toward AI-supported and AI-focused security, including explainable AI, adversarial machine learning, synthetic data, privacy-preserving learning, post-quantum cryptography, critical infrastructures, LLMs, and generative AI security. At the same time, Figure 2 highlights a persistent backbone that cuts across all periods: cryptography, authentication, access control, intrusion detection, and network security remained recurrent pillars even as new platforms, methods, and threats reshaped the research agenda.

4. Community: Who builds SBSeg?

The longitudinal distribution of SBSeg publications shows a gradual but non-linear expansion over its twenty-five editions. As shown in Figure 3, the early years correspond to a smaller and more specialized venue, followed by moderate growth and oscillations during the consolidation period. From the 2010s onward, the three-year moving average indicates a more sustained trajectory, culminating in a sharp increase after 2023 and the highest publication volume in 2025. This recent acceleration suggests that SBSeg has entered a new stage of scale and visibility, consistent with the growing centrality of cybersecurity in digital transformation, privacy regulation, artificial intelligence, critical infrastructures, and socio-technical systems.

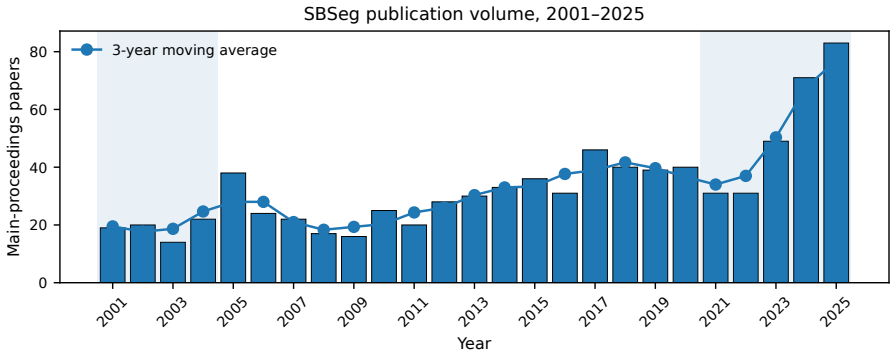

Figure 3. Annual number of main-track SBSeg papers from 2001 to 2025, with a three-year moving average and shaded analytical periods.

4.1. Institutional Footprint

The institutional analysis processed 824 article rows from the SBSeg main track, covering all editions from 2001 to 2025. After institution-name normalization and merging, the corpus comprises 244 unique institutions and 1,311 article–institution incidences after per-row deduplication. The counting rule assigns at most one count to each institution per article row, preventing repeated mentions of the same affiliation within a paper from inflating institutional participation. Figure 4 focuses on the institutions with at least ten incidences, providing a longitudinal view of the main institutional actors in SBSeg.

Figure 4 shows a clear concentration of participation around a small group of recurring institutions. UNICAMP is the most prominent case, with 118 occurrences and sustained activity throughout almost the entire period, including peak years in 2016 and

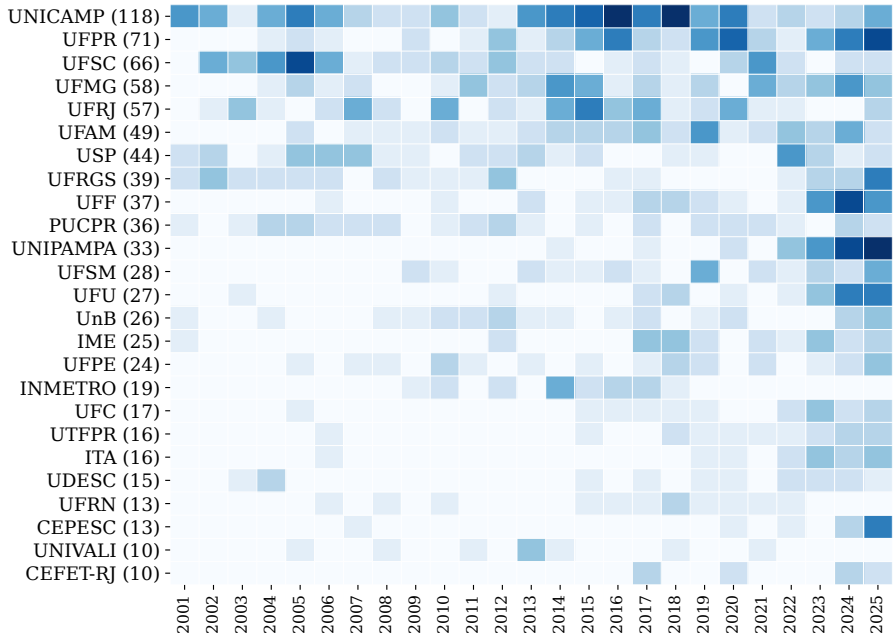

Figure 4. Institutional participation in SBSeg from 2001 to 2025 (>10 papers).

2018, with ten occurrences each. UFPR, UFSC, UFMG, UFRJ, UFAM, USP, UFRGS, UFF, and PUCPR also form a quite visible institutional core. Some of these institutions are associated with specific moments of high intensity, such as UFSC in 2005, UFPR in the late 2010s and especially in 2025, UFRJ around 2015, and UFAM in 2019. These peaks suggest that the institutional presence of SBSeg combines long-term continuity bursts of more intense participation in specific editions.

The most recent editions reveal a second important pattern: the strengthening of institutions whose participation becomes particularly visible after the 2020s. UNIPAMPA is the clearest example, increasing from sparse participation before 2020 to a sequence of strong results between 2022 and 2025, reaching ten incidences in 2025. UFF also shows a strong recent peak, especially in 2023–2025, with its highest value in 2024. UFU, UFRGS, ITA, UFC, UTFPR, and CEPESC likewise become more visible in the final years, with CEPESC reaching its strongest participation in 2025. Overall, the heatmap suggests that SBSeg's institutional community is shaped by both historically persistent centers and more recent waves of institutional consolidation, indicating a broader and increasingly diversified national cybersecurity research community.

4.2. Authorship Footprint

The authorship analysis was conducted at the article level for the SBSeg main track from 2001 to 2025. After author-name normalization and merging, the corpus comprises 1,652 unique authors and 3,026 author–article incidences. The counting rule assigns at most one count to each author per article row, preventing repeated mentions or name variants from inflating individual participation. Figure 5 presents the fifteen most recurrent authors in the corpus, with bars stacked by publication year to show not only total production, but also the temporal distribution of each author's participation.

Figure 5 reveals a small group of authors with recurrent participation across multiple SBSeg editions. Paulo Lício de Geus appears as the most frequent author, with 37

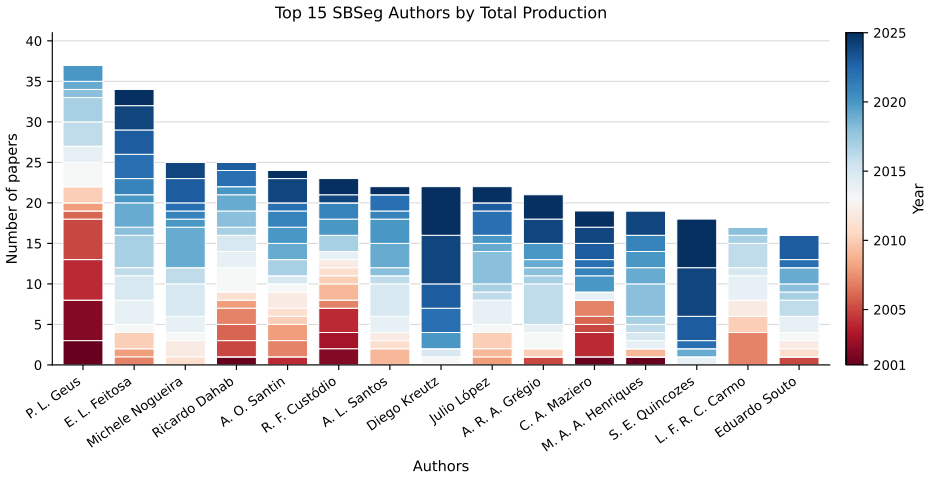

Figure 5. Top 15 authors by SBSeg main-track production from 2001 to 2025.

article incidences and a strong presence in the early and mid-2000s. Eduardo Luzeiro Feitosa follows with 34 incidences and the broadest temporal continuity among the leading authors, remaining active from 2007 to 2025. Ricardo Dahab, Michele Nogueira, Altair Olivo Santin, Ricardo Felipe Custódio, Aldri Luiz dos Santos, Julio López, André Ricardo Abed Grégio, Leonardo Barbosa Oliveira, Carlos Alberto Maziero, and Marco Aurélio Amaral Henriques also illustrate long-term engagement, combining multi-year continuity with specific peaks, such as Michele Nogueira in 2019, André Ricardo Abed Grégio in 2016, Julio López and Marco Aurélio Amaral Henriques in 2018, and Aldri Luiz dos Santos in 2015.

The darker segments in the stacked bars highlight a more recent wave of authorship concentration. Diego Kreutz and Silvio Ereno Quincozes (*e.g.,* coauthors of awarded papers such as [Dresch et al. 2024], [Scherer et al. 2024a], and [Scherer et al. 2024b], among others) stand out in this regard: both reach six incidences in 2024 and again in 2025, indicating strong recent participation despite having shorter overall trajectories than some long-standing authors. Eduardo Luzeiro Feitosa also remains visible in the most recent editions, reinforcing a pattern of continuity across different phases of the symposium. Overall, the authorship distribution suggests that SBSeg's community has been shaped by a combination of early foundational contributors, long-term recurrent authors, and newer high-intensity participants whose activity becomes especially prominent in the 2020s.

4.3. Collaboration Structure

To examine how SBSeg authors relate to one another, we model co-authorship as an undirected weighted graph. Each node is a canonical author, obtained after a name-disambiguation procedure, and an edge links two authors who co-signed at least one main-track paper, with weight equal to the number of papers they share. To keep the figure legible and to center the analysis on authors with a sustained presence in the venue, we restrict the graph to authors with at least three SBSeg papers, which yields 216 authors. From this graph we extract the giant component, detect communities with the Louvain method [Blondel et al. 2008], and compute a force-directed layout. Figure 6 shows the result, with node size encoding productivity and node color encoding community membership.

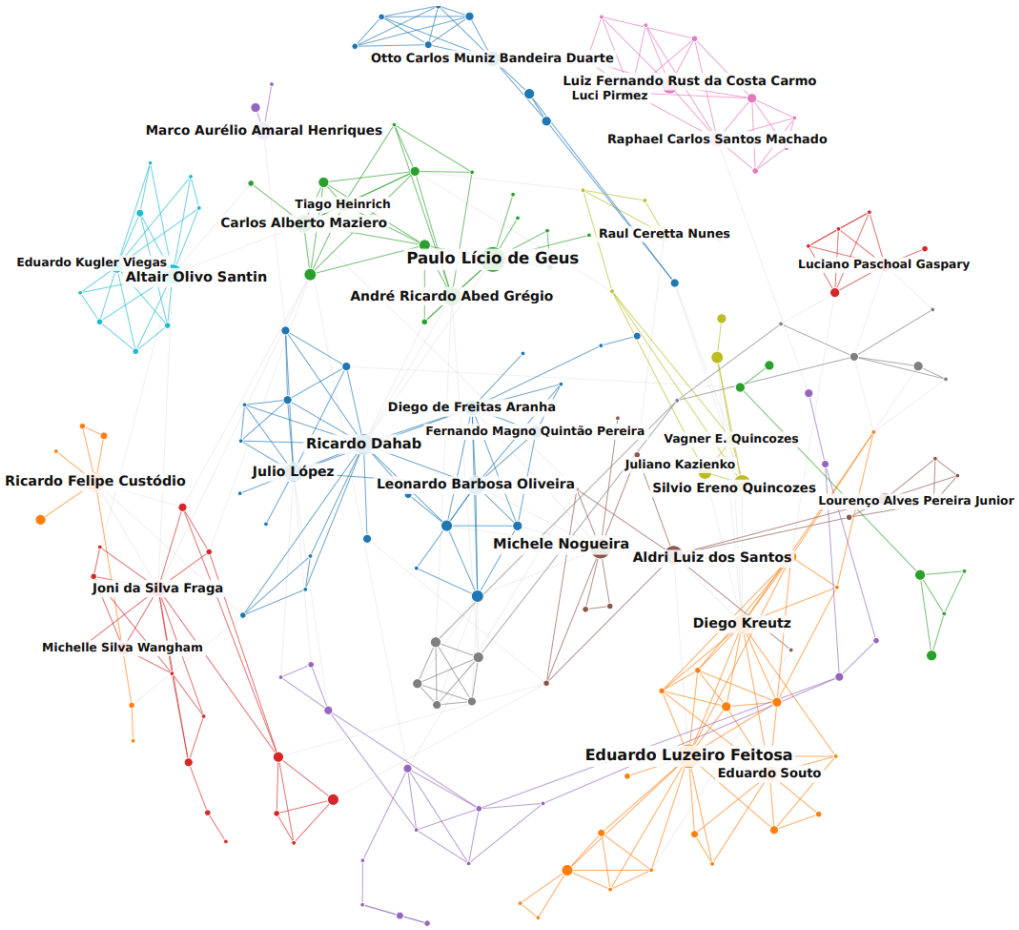

Figure 6. Co-authorship graph of authors with at least three SBSeg papers.

Two structural patterns stand out. First, the SBSeg community is cohesive. The giant component holds 184 of the 216 established authors (≈85%), and even when all 1,479 authors are considered, ≈70% still fall into a single component. SBSeg therefore behaves as a connected collaboration fabric, not as an archipelago of isolated groups. Second, this fabric is polycentric. The Louvain method partitions the giant component into 15 communities of comparable size, with no single dominant cluster. These communities approximate, but do not coincide with, institutional and thematic research groups, so we read them as collaboration neighborhoods rather than as exact laboratory boundaries.

The network is held together by a small backbone of authors, although its membership depends on how centrality is measured. Measured by degree, the most broadly collaborative authors are Eduardo Luzeiro Feitosa, Ricardo Dahab, Diego Kreutz, Altair Olivo Santin, Julio López, and Michele Nogueira, each connected to many co-authors, often concentrated within their own community. Measured by betweenness centrality, which captures how often an author lies on the shortest path between two others, the ranking changes: Aldri Luiz dos Santos is the dominant articulation point, followed by Ricardo Dahab, Carlos Alberto Maziero, Michele Nogueira, and Diego Kreutz.

The two rankings share only four of their top ten authors, and some effective bridges have a modest number of collaborators yet connect three or four otherwise-separated communities, as in the case of Maziero. In the evaluated graph, the most

prolific collaborators are thus not necessarily the ones that hold the venue together: part of SBSeg's cohesion rests on boundary-spanning authors who link communities they do not dominate. Because Figure 6 is restricted to established authors and to the giant component, it deliberately understates the many one-time collaborations that surround this core.

5. Open Challenges and Opportunities

The longitudinal analysis suggests that SBSeg is well positioned to address three emerging research fronts that extend its historical strengths in cryptography, network security, intrusion detection, privacy, and AI-supported defense. These fronts also help answer RQ3 by indicating how the venue's trajectory points to the future of cybersecurity research in Brazil.

Agentic AI security. Recent SBSeg editions already include work on LLMs, jailbreaks, guardrails, data-extraction attacks, and AI-supported defense. A natural next step is agentic AI security, where autonomous agents interact with tools, APIs, code, files, and external systems. In this setting, prompt injection, tool misuse, sandboxing, inter-agent trust, and accountability become central open problems.

Post-quantum transition. Cryptography has been a persistent SBSeg backbone, and recent editions show growing attention to post-quantum schemes and constrained environments. The next challenge is not only designing primitives, but supporting migration at scale through hybrid modes, crypto-agility, public-key infrastructure updates, and protection against harvest-now-decrypt-later threats.

AI-native communication infrastructures. SBSeg has also begun to address programmable and next-generation infrastructures such as SDN, Open RAN, 6G, Kubernetes, IoT, and smart grids. As these systems become increasingly software-defined and AI-operated, new challenges emerge around adversarial learning, model evasion, multi-vendor trust, in-network detection, and resilient operation of critical communication infrastructures.

6. Conclusion

This paper presented an article-level longitudinal review of 25 years of SBSeg main-track research, covering 25 proceedings volumes and 824 papers. Regarding RQ1, the results show that SBSeg preserved foundational areas such as cryptography, authentication, access control, intrusion detection, and network security, while progressively incorporating privacy, IoT, cloud computing, SDN, blockchain, post-quantum cryptography, machine learning, adversarial AI, LLMs, and generative AI security.

Regarding RQ2, SBSeg evolved into a broader and more cohesive national research community, blending historically recurrent authors and institutions with recent institutional consolidation, high-intensity contributors, and a polycentric collaboration structure, evidence that its growth was quantitative as well as social and institutional. Regarding RQ3, the co-evolution of topics and community reveals SBSeg's maturation from information and computer systems security into a modern, AI-native cybersecurity forum, pointing to open fronts such as agentic AI security, post-quantum transition, and AI-native communication infrastructures. Threats related to metadata completeness, ter-

minology variation, and interpretive classification were mitigated through normalization, manual validation, and LLM-assisted consistency checking under author supervision.

References

Afonso, V. M., de Amorim, M. F., Ellery, E., Grégio, A. R. A., Junquera, G. B., Schick, G. A. K., Dahab, R., and de Geus, P. L. (2013). Um sistema para análise e detecção de aplicações maliciosas de android. In *Anais do Simpósio Brasileiro de Segurança da Informação e de Sistemas Computacionais*.

Bezerra, V. H., da Costa, V. G. T., Martins, R. A., Junior, S. B., Miani, R. S., and Zarpelão, B. B. (2018). Providing iot host-based datasets for intrusion detection research. In *Anais do Simpósio Brasileiro de Segurança da Informação e de Sistemas Computacionais*.

Blondel, V. D., Guillaume, J.-L., Lambiotte, R., and Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment*, 2008(10):P10008.

Botacin, M., da Rocha, V. F., de Geus, P. L., and Grégio, A. (2017). Analysis, anti-analysis, anti-anti-analysis: An overview of the evasive malware scenario. In *Anais do Simpósio Brasileiro de Segurança da Informação e de Sistemas Computacionais*.

Bragança, H., Rocha, V., Souto, E., Kreutz, D., and Feitosa, E. (2023). Explaining the effectiveness of machine learning in malware detection: Insights from explainable ai. In *Anais do XXIII Simpósio Brasileiro de Segurança da Informação e de Sistemas Computacionais*. Sociedade Brasileira de Computação.

Brandão, J. E. M. S., da Silva Fraga, J., and Mafra, P. M. (2005). Composição de idss usando web services. In *Anais do Simpósio Brasileiro de Segurança da Informação e de Sistemas Computacionais*.

Campello, R. S., Weber, R. F., da Silveira Serafim, V., and Ribeiro, V. G. (2001). O sistema de detecção de intrusão asgaard. In *Anais do Workshop em Segurança de Sistemas Computacionais*.

Ceron, J. M., Granville, L., and Tarouco, L. (2009). Taxonomia de malwares: Uma avaliação dos malwares automaticamente propagados na rede. In *Anais do Simpósio Brasileiro de Segurança da Informação e de Sistemas Computacionais*.

Cervantes, C., Poplade, D., Nogueira, M., and Santos, A. (2014). Um sistema de detecção de ataques sinkhole sobre 6lowpan para internet das coisas. In *Anais do Simpósio Brasileiro de Segurança da Informação e de Sistemas Computacionais*.

Corrêa, J. L., Proto, A., and Cansian, A. M. (2008). Modelo de armazenamento de fluxos de rede para análises de tráfego e de segurança. In *Anais do Simpósio Brasileiro de Segurança da Informação e de Sistemas Computacionais*.

de Albuquerque, J. P., Isenberg, H., Krumm, H., and de Geus, P. L. (2005). Gerenciamento baseado em modelos da configuração de sistemas de segurança em redes de larga escala. In *Anais do Simpósio Brasileiro de Segurança da Informação e de Sistemas Computacionais*.

de Camargo, L. F., Reis, C., Paiola, P. H., Papa, J. P., Brega, J. R. F., and da Costa, K. A. P. (2021). Métodos de aprendizado de máquina adversariais na detecção de anomalias

em redes de computadores. In *Anais do XXI Simpósio Brasileiro de Segurança da Informação e de Sistemas Computacionais*. Sociedade Brasileira de Computação.

Dresch, F. N., Scherer, F. H., Quincozes, S. E., and Kreutz, D. (2024). Modelos interpretáveis com inteligência artificial explicável (xai) na detecção de intrusões em redes intra-veiculares controller area network (can). In *Simpósio Brasileiro de Cibersegurança (SBSeg)*, pages 445–460. SBC.

Ferreira, A. S. and de Geus, P. L. (2013). Uma arquitetura para monitoramento e detecção de anomalias de segurança para nuvens computacionais. In *Anais do Simpósio Brasileiro de Segurança da Informação e de Sistemas Computacionais*.

Henke, M., Costa, C., dos Santos, E. M., and Souto, E. (2011). Detecção de intrusos usando conjunto de k-nn gerado por subespaços aleatórios. In *Anais do Simpósio Brasileiro de Segurança da Informação e de Sistemas Computacionais*.

Laureano, M., Maziero, C., and Jamhour, E. (2004). Proteção de detectores de intrusão através de máquinas virtuais. In *Anais do Workshop em Segurança de Sistemas Computacionais*.

Mazetto, B. O. and Zarpelão, B. B. (2024). A triad of defenses to mitigate poisoning attacks in federated learning. In *Anais do XXIV Simpósio Brasileiro de Segurança da Informação e de Sistemas Computacionais*. Sociedade Brasileira de Computação.

Motta, G. H. M. B. and Furuie, S. S. (2002). Um modelo de autorização contextual para o controle de acesso baseado em papéis. In *Anais do Workshop em Segurança de Sistemas Computacionais*.

Neto, A. L. M., Nogueira, M., Patil, H. K., Ítalo Cunha, Loureiro, A. A. F., and Oliveira, L. B. (2015). Criptografia baseada em identidade: Uma análise comparativa sob a perspectiva da internet das coisas. In *Anais do Simpósio Brasileiro de Segurança da Informação e de Sistemas Computacionais*.

Oliveira, A. K. D. S. and López, J. (2013). Implementação em software do esquema de assinatura digital de merkle e suas variantes. In *Anais do Simpósio Brasileiro de Segurança da Informação e de Sistemas Computacionais*.

Oliveira, L. B., Wong, H. C., Bern, M., Habib, E., Loureiro, A. A. F., and Dahab, R. (2005). Secleach – a random key distribution solution for securing clustered sensor networks. In *Anais do Simpósio Brasileiro de Segurança da Informação e de Sistemas Computacionais*.

Paim, K. O., Mansilha, R. B., Kreutz, D., Franco, M. F., and Cordeiro, W. (2025). Exploiting latent space discontinuities for building universal llm jailbreaks and data extraction attacks. In *Anais do XXV Simpósio Brasileiro de Cibersegurança*. Sociedade Brasileira de Computação.

Peres, A. and Weber, R. F. (2003). Considerações sobre segurança em redes sem fio. In *Anais do Workshop em Segurança de Sistemas Computacionais*.

Piedrahita, A. F. M., Rueda, S., Mattos, D. M. F., and Duarte, O. C. M. B. (2015). Flowfence: Um sistema de defesa contra ataques de negação de serviço para redes definidas por software. In *Anais do Simpósio Brasileiro de Segurança da Informação e de Sistemas Computacionais*.

Ribeiro, P. S. and Custódio, R. F. (2002). Um protocolo criptográfico para comunicação anônima segura em grupo. In *Anais do Workshop em Segurança de Sistemas Computacionais*.

Scherer, F. H., Dresch, F. N., Quincozes, S. E., Kreutz, D., and Quincozes, V. E. (2024a). Iwshap: Um método de seleção incremental de características para redes can baseado em inteligência artificial explicável (xai). In *Simpósio Brasileiro de Cibersegurança (SBSeg)*, pages 351–366. SBC.

Scherer, F. H., Dresch, F. N., Quincozes, S. E., Kreutz, D., and Quincozes, V. E. (2024b). Iwshap: Uma ferramenta para seleção incremental de características utilizando iwss e shap. In *Simpósio Brasileiro de Cibersegurança (SBSeg)*, pages 105–112. SBC.

Silva, E. S., Santin, A. O., Jamhour, E., Maziero, C. A., and Toktar, E. (2007). Extensões ao modelo rbac de restrições para suportar obrigações do uconabc. In *Anais do Simpósio Brasileiro de Segurança da Informação e de Sistemas Computacionais*.

Vasquez, G., Miani, R. S., and Zarpelão, B. B. (2017). Flow-based intrusion detection for scada networks using supervised learning. In *Anais do Simpósio Brasileiro de Segurança da Informação e de Sistemas Computacionais*.