Mapping the Cloud: A Mixed-Methods Study of Cloud Security and Privacy Configuration Challenges

Sumair Ijaz Hashmi*[†][‡][||], Shafay Kashif*[§][||], Lea Gröber[¶][||], Katharina Krombholz[†] and Mobin Javed[||]
[†]CISPA Helmholtz Center for Information Security, Germany
[‡]Saarland University, Germany
[§]The University of Auckland, New Zealand
[¶]International Computer Science Institute (ICSI), USA
[||]Lahore University of Management Sciences (LUMS), Pakistan

***Abstract*—Misconfigurations in cloud services remain a leading cause of security and privacy incidents, often stemming from the complexity of configuring cloud platforms. To better understand these challenges, we analyzed approximately 251,900 security- and privacy-related Stack Overflow posts spanning from 2008 to 2024. Using topic modeling and qualitative analysis, we systematically mapped cloud use cases to their associated security and privacy configuration challenges, revealing a comprehensive landscape of the hurdles cloud operators faced. We identified both technical and human-centric issues, including problems related to insufficient documentation and the lack of context-aware tooling tailored to operators' environments. Notably, authentication and access control challenges appeared in all identified use cases, cutting across nearly every stage of cloud deployment, integration, and maintenance. Our findings underscore the need for usable, tailored, and context-sensitive support tools and resources to help developers securely configure cloud services.**

I. Introduction

Cloud computing platforms such as Amazon Web Services (AWS) [1], Google Cloud Platform (GCP) [2], and Microsoft Azure [3] have become integral to modern software development due to their scalability, affordability, and ease of deployment [4]. Recent industry reports show that over half of organizations migrated workloads to the cloud in 2024, and 79% of large enterprises now use multiple cloud providers [5], [6]. Yet, as cloud adoption accelerates, so do its security and privacy risks. Despite their apparent convenience, securely configuring cloud services remains a significant challenge for various professionals involved in cloud operations, such as developers, system administrators, and network engineers, whom we collectively refer to as *cloud operators* in this paper. Misconfigurations have repeatedly been identified as the leading cause of cloud security breaches [6], [7], [8]. In July 2024, for example, attackers accessed call and text records of over 100 million AT&T customers by exploiting improperly secured Snowflake databases that lacked enforced multi-factor authentication [9]. Similarly, a misconfiguration in Toyota's cloud infrastructure exposed vehicle and customer data for over a decade [10], [11]. These are not one-off incidents, but are rather systemic issues in the cloud ecosystem; a recent report investigating production assets on major cloud providers such as Azure, AWS, and Google Cloud found cloud assets have 115 vulnerabilities on average [12], [13]. This landscape underscores the urgent need to understand where and how such misconfigurations persist, from a human-centric perspective.

Although researchers have conducted studies on technical vulnerabilities in the cloud infrastructure [14], [15] and characterized security misconfigurations from the broad perspective of administrators [16], the specific challenges faced by everyday cloud operators remain largely underexplored—aside from two investigations into the self-hosters' security mindsets and challenges [17], [18]. We bridge this gap by presenting the first large-scale empirical mapping of operators' real-world cloud use cases to security and privacy configuration hurdles they face, laying foundations for improving cloud security posture.

We leverage Stack Overflow, a widely used platform where cloud operators seek and share practical guidance. Prior research has shown that such informal information sources shape how security is understood and implemented [19], [20], and that these sources now also feed large language models (LLMs) like ChatGPT [21], [22], which are increasingly integrated into developers and administrators' workflows [23]. Analyzing operator interactions on Stack Overflow thus provides a unique lens into the real-world challenges of cloud security configuration, and how the cloud operator community responds to them. Crucially, we go beyond raw vulnerability counts to uncover human-centric usability issues, such as inadequate documentation, poor tooling, and knowledge gaps, that persist in cloud security workflows.

We conducted a mixed-methods analysis to examine the use cases, security and privacy challenges, and support mechanisms cloud operators rely on when configuring security and privacy in the cloud. We also focus on accepted answers to

* Both authors contributed equally to this research.

Network and Distributed System Security (NDSS) Symposium 2026
23-27 February 2026, San Diego, CA, USA
ISBN 979-8-9919276-8-0
https://dx.doi.org/10.14722/ndss.2026.241302
www.ndss-symposium.org

identify widely endorsed solutions and advice. Specifically, we ask:

- **RQ1:** What security and privacy-related configuration challenges do cloud operators face across the cloud ecosystem?
- **RQ2:** What types of human-centric challenges are associated with these configuration tasks?
- **RQ3:** What solution strategies do accepted answers provide for addressing these challenges?

To answer these questions, we collected approximately 251,900 cloud-related security and privacy posts from Stack Overflow, spanning from August 2008 to March 2024. We applied topic modeling to uncover major themes and then conducted a qualitative analysis of a stratified sample of 625 posts. From this analysis, we identified a range of security and privacy–related use cases and configuration challenges encountered by cloud operators. By mapping use cases to associated challenges, we constructed a landscape of the most pressing hurdles faced during cloud security configuration. We found that configuration issues commonly revolve around: (i) authentication and access control, (ii) secure communication and encryption, (iii) network configuration and management, (iv) logging and monitoring, and (v) database backup and recovery. Among these, authentication and access control emerged as the most pervasive challenge, cutting across nearly all cloud use cases. We also observed that many challenges stem from impractical, incomplete, or overly generic documentation, as well as a lack of context-aware tools to guide users through complex configuration tasks. Our findings highlight the importance of designing more usable advice resources, documentation, and developer support tools to mitigate misconfigurations. We discuss the potential of AI-powered assistants and other interactive interfaces to guide users through security configurations.

To support open science, we provide a full replication package, including data, analysis scripts, and our codebook[1].

II. Related Work

We review work on (i) security and privacy issues in the cloud, (ii) studies focusing on human factors in the cloud, and (iii) studies leveraging online developer forums.

A. Security and Privacy in the Cloud

Studies highlight the risks cloud platforms pose to consumer privacy [24], [25], [26], [27]. Similarly, security in the cloud is a multifaceted challenge, and several studies highlight security issues, such as those associated with data integrity and storage, availability of cloud services, and network security [28], [29], [30], [31]. Researchers have explored these issues through various lenses, such as by surveying high-level cloud provider and tenant issues [15], from the perspective of the architecture frameworks and cloud technologies [32], or by analyzing real-life deployment of cloud services for issues in backend applications [14]. Alrawi et al. in particular highlight vulnerabilities related to unpatched or legacy operating systems, poor coding practices in cloud, and misconfigured SSL configurations for communication [14]. We complement this line of work by focusing on operators' configuration tasks and underpinning issues that precede many of the vulnerabilities identified by these studies. Rather than solely proposing new technical countermeasures, we map where and how configuration hurdles emerge in practice. Although prior research has explored technical solutions to specific security and privacy problems in the cloud [33], [34], [35], [36], [37], [38], [39], [40], [41], [42], [43], such as proposing stateful least privilege authorization tokens for cloud services to prevent unauthorized access [42], comparatively less is known about which configuration steps cloud operators struggle with. Our study helps empirically locate these difficulties. We address this gap by uncovering cloud operators' goals and challenges when implementing and managing cloud security and privacy.

B. Studies on Human Factors

Studies have shown that end-users often have incomplete and incorrect mental models of file storage, sharing, and deletion in commercial clouds [44], [45], [46], [47], [48], [49]. Studies with expert users were more narrow in scope and investigated: (i) misconceptions and challenges faced by network and server administrators [50], [51], [16], [52]; and (ii) usability challenges in configuring transport layer security [53], [54], [55], [56].

Two studies by Gröber et al. focus on individuals choosing to self-host their cloud infrastructure; the first investigating their operation mechanisms, threat models, and defense mechanisms [17], and the second investigating their motivations, such as an interest in technology, hosting skills, and a "maker" self-identity [18]. They find that despite their motivation, self-hosters are limited in their technical knowledge and rely on others to help configure their desired use cases [17], [18]. However, while these studies examine security mindsets of niche operator groups (e.g., self-hosters) [17], [18]; large-scale, cloud-wide mappings of operator challenges are lacking. Our analysis of problem-driven experiences complements prior work, and captures a broader and more ecologically valid view of real-world challenges as they emerge in practice, across multiple cloud providers and operator roles, without the constraints of recruitment or recall bias.

We review research on documentation, security education, and security culture in organizations as these may shape how operators overcome security misconfigurations:

Documentation: In a study on documentation page-view logs of four Google Cloud services, researchers found that different patterns of documentation usage emerge based on several user characteristics, such as users' background, context, and prior experience with the cloud service, among others [57]. These researchers recommend making documentation pages more usable by providing personalized guidance based on user needs and background [57]. Prior work has also recommended making documentation pages more interactive to improve usability [58]. Regarding security, prior work has

[1] View Zenodo Repository here.

recommended re-ranking Google search results on security-related programming queries based on secure coding metrics [59].

Security Education: While there has been some pedagogical work in cloud computing [60], system administrators generally do not credit formal education as the primary source of their expertise; instead, they emphasize experiential, "hands-on" learning [61]. An exploratory study on self-hosters with varying levels of expertise found no clear relation between professional education and structured approaches to security [17].

Security Culture: Gröber et al. characterized the social embeddedness of cloud operators and found that their approaches to social support impacted their security strategies [17]. Weir et al. [62] and Ryan et al. [63] highlight how developers' security culture and awareness influences secure-coding practices. Ryan et al., through a survey of 1,100 developers, show that developer teams tend to perform security activities for compliance purposes but still lack a strong security culture in their organizations [63]. Similarly, Weir et al. show that such software teams can improve their security practices through facilitated workshops that improve on security posture [62].

We extend prior human-centered studies by conducting a bottom-up analysis of security and privacy–related discussions among a broad set of cloud operators (not only specific subpopulations), grounding our categories in the issues they raise on Stack Overflow. Our findings both reinforce usability and mental-model issues observed in narrower system administration contexts and reveal how these problems manifest across the broader cloud pipeline.

C. Studies Focusing on Developer Forums

Online developer forums are a community of practice where developers informally discuss programming, set-up, and configuration related issues and learn from one another [64]. They serve as an essential source of knowledge for developers [65], [66], [67] and therefore, provide a window into the challenges developers face. Stack Overflow is one such platform designed for developers, encompassing various topics such as website development, databases, version control, and security [66]. It serves as a valuable data source regarding developers—who are often difficult to engage for empirical research [68], [69], [70]. Further, despite the emerging rise of LLMs such as ChatGPT as a replacement of Stack Overflow in software development, a recent study shows that users still prefer human answers from Stack Overflow as they find human answers to be more correct, concise, and useful than LLM-generated content [71]. This forum has been widely studied by researchers to identify prevalent security and privacy topics, as well as the challenges in meeting security and privacy requirements [72], [64], [73], [74], [75], [76]. Stack Overflow is an important resource to systematically study as developers use security and privacy-related code snippets from Stack Overflow and do not update these code snippets despite the snippets having received bug or security fixes on the developer forum [77]. However, such prior analyses of Stack Overflow security and privacy content considered non-cloud domains (e.g., operating systems, social networks) or specific vulnerability classes [72], [64], [73], [74], [75], [76]. In contrast, we focus specifically on the security and privacy configuration in cloud contexts, and we map use cases to challenges at scale.

III. Methodology

We conducted a three-phase mixed-methods study to analyze security and privacy-related challenges in cloud service configuration on Stack Overflow, using data from 2008 (the year StackOverflow was launched) to 2024. The study comprised three phases: (i) data collection, (ii) topic modeling using Latent Dirichlet Allocation (LDA), and (iii) qualitative analysis through bottom-up thematic coding. Finally, we estimated the potential spread of themes across the full dataset via a conservative, topic-based upper-bound mapping. Figure 1 provides an overview of our methodology.

To capture long-term trends and validate the relevance of our findings over time, we split the dataset into two parts, following guidance from Jallow et al. [78] on temporal validation in Stack Overflow studies. We applied all three analysis phases to each part: a larger dataset spanning August 2008 to 25th Sept. 2022, and a more recent dataset covering 26th Sept. 2022 to March 2024. We split the dataset on Sept. 2022 to account for the introduction of large language models (LLMs) to the public in November 2022.

A. Phase 1: Data Collection

We collected cloud-related security and privacy posts from Stack Overflow using a multi-step process described below.

1) Tags Collection and Obtaining an Initial Cloud Dataset: To define the scope of cloud computing, we focused on major hosting providers that offer cloud services. We refer to all help-seekers, e.g., developers, system administrators, self-hosters, and network engineers, collectively as *cloud operators*.

As a starting point, we used the Stack Overflow Developer Survey [79], an annual survey of the developer community with over 89,000 participants in 2023 alone [80]. Since 2020, the survey has included questions about popular cloud platforms. By aggregating responses from the 2020–2023 surveys, we identified 23 commonly used hosting providers—including Amazon AWS, Microsoft Azure, Google Cloud Platform, Firebase, Heroku, and others (see Table III for the full list). Using this list, we generated wildcard strings to retrieve all Stack Overflow posts tagged with provider names (e.g., `amazon`) or containing the keyword `cloud` in their tags. Including the generic `cloud` tag ensured broader coverage, capturing relevant posts outside the top providers—such as discussions on `nextcloud`. Then, we extracted all posts that contained at least one such tag. We queried the publicly available Stack Overflow dataset on Google BigQuery [81] and extracted 815,272 posts dated from Aug. 2008 to Sept. 2022. Using the Stack Exchange Data Dump [82] on the Internet Archive [83] we extracted 117,608 posts dated from Sept. 2022 to March 2024. We used a different data source here as the Google Big Query dataset only had posts up until 25.09.2022.

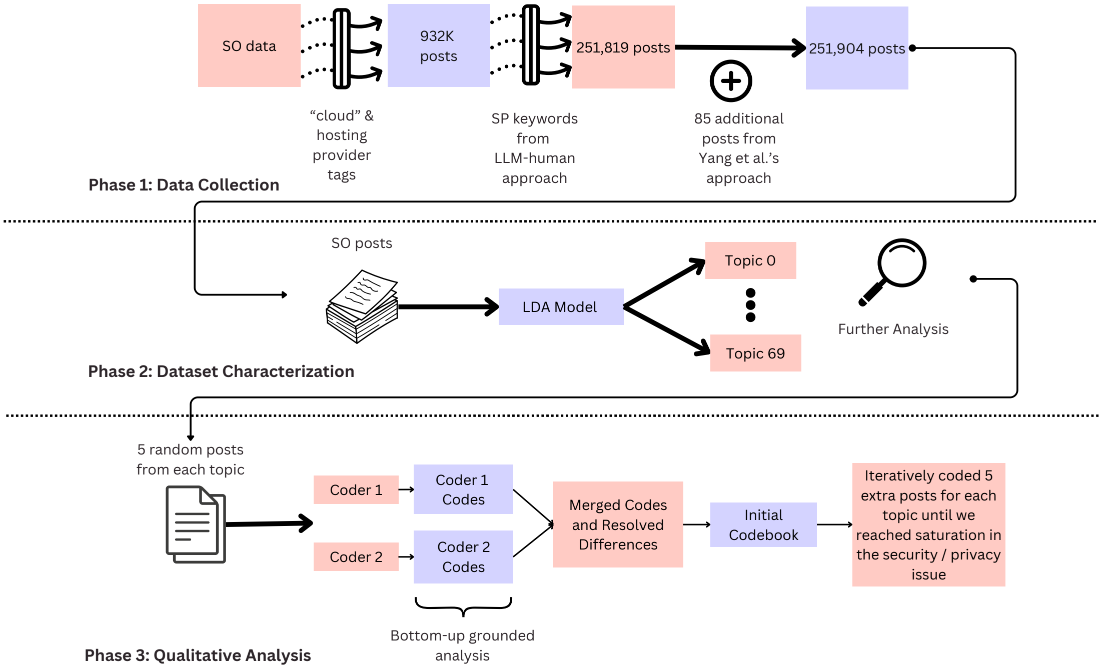
SO data
932K posts
251,819 posts
251,904 posts
"cloud" & hosting provider tags
SP keywords from LLM-human approach
85 additional posts from Yang et al.'s approach
Phase 1: Data Collection
SO posts
LDA Model
Topic 0
Topic 69
Further Analysis
Phase 2: Dataset Characterization
5 random posts from each topic
Coder 1
Coder 2
Coder 1 Codes
Coder 2 Codes
Merged Codes and Resolved Differences
Initial Codebook
Iteratively coded 5 extra posts for each topic until we reached saturation in the security / privacy issue
Bottom-up grounded analysis
Phase 3: Qualitative Analysis

Fig. 1: Flowchart illustrating our data collection and analysis methodology.

2) Filtering Security- and Privacy-Related Posts: We then filtered the cloud dataset to extract posts related to security or privacy. To do this, we used a keyword-based filtering strategy guided by large language model output. Prior work has investigated how LLMs can be used to facilitate scientific research by subjecting it to human-in-the-loop feedback and review [84], [85], [86], [87], [88]. We implemented a similar human-in-the-loop use of GPT4 [89] to generate a list of security- or privacy-related keywords, which we then manually analyzed and reviewed in an iterative process to obtain a diverse set of relevant posts while minimizing false positives.

a) Selection of Security and Privacy Keywords: We used GPT-4 [89] to generate an initial list of 450 keywords relevant to security and privacy. Our prompts (included in Table IV, Appendix B) were profession-agnostic and based on NIST standards related to cloud computing [90], [91] including use cases, threat models, and configuration practices. Two authors independently formulated prompts, which were discussed and merged before querying the LLM. After five prompts, we observed keyword saturation. We manually reviewed the resulting keyword list through a two-stage process. First, for each keyword, we obtained a random sample of 25 matching posts from the cloud dataset. Posts were independently labeled by two researchers to be security or privacy related or not. This approach allowed us to identify a set of "gray" keywords that produced a majority of false positives in the reviewed sample of 25, as the content of the posts was not related to security or privacy, and the keyword was being used in another context or meaning. Removing "gray" keywords inevitably risks excluding some relevant posts; however, in most such cases other security and privacy keywords still captured the posts. An example of such a keyword was ”updates,” which we initially incorporated to extract posts related to maintaining, updating, and patching cloud components, but we later removed it as the term was mostly used in other contexts, such as updating data in databases. Next, the entire team discussed "gray" keywords based on the labeled posts and under consideration of the NIST guidelines. After reaching consensus, 17 keywords were removed, leaving 433 keywords. In numbers: these ”gray” keywords yielded 104,625 posts out of which 54,243 (51.8%) are already captured by our dataset. Our earlier manual review of posts containing the keywords mentioned along with a recheck of a sample of posts not captured by other keywords, showed that these posts were not relevant to our research questions. Including them would risk reducing the validity of our dataset.

b) Extraction of Security and Privacy Related Posts: We used the resulting keywords to extract all posts that mentioned any of them in the title, body, or tags. This yielded 214,133 security- and privacy-related posts from the inital 815,272 for the 2008-25.09.2022 dataset and 37,686 security- and privacy-related posts from the initial 117,608 for the 26.09.2022-2024 dataset.

3) Validation Against Prior Research: To validate our filtering approach, we replicated Yang et al.'s method [75] on the 2008–2022 dataset. Their approach identifies security-related posts by analyzing the co-occurrence of tags with `security` and then applying thresholding metrics to shortlist security-related tags. Posts are subsequently extracted based on the resulting tags, a method that has been adopted in follow-up studies [92], [93]. We extended this methodology to include `privacy`-related tags in an analogous manner.

We experimented with different threshold values to maximize recall and found that the original thresholds from Yang et al. [75] missed key tags such as `authentication`. After manual testing, we set the thresholds to $Thre1 = 0.1$ and $Thre2 = 0.0005$. Using this modified method, we retrieved 92,888 security- or privacy-related posts. However, when we applied our earlier cloud tag filters to this subset, only 3,340 posts remained, substantially fewer than the 214,133 posts retrieved through our LLM-driven approach.

To understand why our approach yielded substantially more posts, we closely examined the tags shortlisted by Yang et al.'s method [75]. We found that their approach failed to capture several key security and privacy topics—such as authentication and authorization—that are well-represented in our dataset. For instance, in the dataset produced using Yang et al.'s method, the tag `authentication` appears alongside `security` fewer than 100 times. We assume this is likely because such tags do not frequently co-occur with `security` on Stack Overflow, as the tagging process is user-driven and discretionary. Question askers and moderators may omit relevant tags, use inconsistent terminology, or simply be unaware of the best tags to apply. This suggests that relying solely on tag co-occurrence–based filtering may produce a limited and incomplete dataset for identifying security and privacy posts. Further, we compared our dataset with the one produced using Yang et al.'s approach by calculating the intersection and set differences between the two. We found that our approach already included 97.5% (3,255 posts) of the posts from their dataset. Only 85 posts from their dataset were missing in ours. We manually reviewed these 85 posts and confirmed that all were indeed relevant to cloud security or privacy. We added them to our final dataset.

Since our method successfully captures the vast majority of posts from Yang et al.'s dataset—while also retrieving many additional relevant posts that their method misses—we argue that our approach is valid and more comprehensive. Although not exhaustive, it provides a substantially broader and richer dataset for analyzing cloud-related security and privacy discussions on Stack Overflow. Based on this evaluation, we consider our method appropriate for this task and therefore did not reapply Yang et al.'s approach for the 2022–2024 dataset. After merging the results from the prior work's approach, our final 2008-2022 dataset contains 214,218 posts.

B. Phase 2: Dataset Characterization

To gain an initial understanding of the datasets' content and themes, we conducted topic modeling. Moreover we analyzed linked help resources to examine how cloud operators seek external guidance. Details of these analyses are below.

1) Topic Modeling: We applied Latent Dirichlet Allocation (LDA) [94] to extract latent topics from the dataset. LDA is an unsupervised topic modeling technique that identifies naturally occurring word clusters without relying on human-assigned labels. It is widely used in prior work on Stack Overflow content analysis [95], [66], [65], [75], [96], [92], making it a suitable choice for our study.

a) Preprocessing Data: Before applying LDA, we performed several preprocessing steps by (i) removing HTML tags and non-alphabetic characters, (ii) converting all text to lowercase, and (iii) eliminating stopwords and specific identifiers like email addresses and URLs. We also applied stemming and condensed the data by merging titles, bodies, and tags into single documents while removing all numeric values and code. Codes and numerical values were removed to generate topic clusters based only on natural language data. This avoids generating topics based on coincidentally frequently occurring words, such as print statements.

b) Model Configuration and Optimization: We used Gensim's implementation of LDA [97] and experimented with topic counts ranging from 15 to 250 (in increments of 5) to identify an optimal configuration. We evaluated each model using two established metrics: (i) topic *coherence*, which measures semantic consistency within topics, and (ii) *cosine similarity*, which captures topic distinctiveness. Both metrics have been used in prior research to determine the number of topics in LDA [93], [98]. To determine the optimal topic count, we used the elbow method [98] on the cosine similarity curve and cross-checked with coherence scores. Based on this analysis, for the 2008-2022 dataset we selected 50 topics—this value produced one of the five highest coherence scores and marked a clear elbow point in the cosine similarity curve (see Figure 4). Analogously, we selected 20 topics for the 2022-2024 dataset (see Figure 5). Two researchers jointly interpreted and annotated the topics with labels by examining the most representative words for each. See the extended supplementary material for an overview of all uncovered topics.

2) Usage of Documentation and Other Help Resources: In addition to analyzing topic content, we examined how cloud operators referenced external resources when posing questions. This provides insight into what types of help operators seek when addressing cloud security and privacy issues. We began by extracting all hyperlinks (`<a>` tags) from each post to identify linked domains. We filtered out non-informative or auto-generated links—such as "localhost," "stackimgur," (used for embedded images), and "example.com." Next, we tallied the frequency of each root domain and manually categorized the top 63 domains (those appearing in at least 100 posts). These were manually grouped into categories e.g., official documentation, support portals, code samples, blogs, forums.

C. Phase 3: Qualitative Analysis

To complement our topic modeling results with a deeper understanding of user concerns and solutions, we conducted a qualitative analysis on stratified samples from each topic. Topics yielded by LDA were not mutually exclusive in terms of use cases and security challenges. Hence, the qualitative analysis allowed us to categorize overarching and recurring themes. Our qualitative coding was inspired by Grounded Theory techniques (open and axial coding) [99], [100], [101], [102], but we did not aim to build a formal theory; we stopped after axial coding to organize themes around our RQs.

1) Coding Process: We applied bottom-up grounded analysis via open coding to obtain insights from the data [102], [99], [100], [101] starting with the 2008-2022 sample. Our goal was to identify patterns in the questions regarding what the question-askers were trying to achieve, what key issues they were facing, what security and privacy issue or mechanism was being discussed, and what strategies were being used in the accepted answers to respond to the questions. Additionally, we explore the context behind using the help resources identified in Section III-B2 by coding why cloud operators referred to these resources and what challenges they expressed to have faced with them, if any.

Two researchers independently coded a stratified random sample of 5 posts per topic based on the themes described above. Then, they discussed their codes and code assignments with each other to resolve disagreements and construct an initial version of the codebook. They followed an iterative process of coding five more posts from each topic until they reached thematic saturation [103], [104] with respect to the security or privacy aspect of the posts for that topic. Sampling was stratified by topic to ensure coverage of the datasets' thematic diversity. For the 2008-2022 dataset, this resulted in 525 posts (184 with accepted answers), and 100 posts for the 2022-2024 dataset.

After open coding, the authors jointly discussed the emerging themes by grouping codes in the codebook and conducting axial coding to identify additional insights. Thereby we identified four axial categories (*use cases*, *configuration challenges*, *human-centric challenges*, and *answer strategies*). The codebook is attached in the extended supplementary material.

2) Relevance Across Time: Post-LLM Analysis: To assess whether identified security and privacy challenges and cloud use cases persisted over time, we applied our final codebook of the 2008-2022 dataset to the 2022–2024 sample. No new high-level themes emerged, indicating that key security and privacy challenges have remained consistent, even in the LLM era. We excluded accepted answer analysis for this recent set, as accepted answers now may also contain LLM-generated content that may distort assessments of human problem-solving strategies, an important direction for future work.

D. Connecting Topic Modeling with the Qualitative Analysis

To assess the prevalence of the challenges identified through qualitative analysis across the entire dataset, we linked them back to the topic modeling results. While our qualitative analysis does not aim for generalizability, this connection allows us to estimate an *upper bound* on the distribution of use cases and configuration challenges across topics. We began by merging posts from the 2008–2022 and 2022–2024 qualitative samples. Each post in this merged sample was annotated with a use case and configuration challenge (as established in Section III-C). Posts sharing the same pair of annotations were grouped into a *cell*, representing a specific (*use case*, *configuration challenge*) combination. To estimate the potential spread of each cell, we identified all topics represented by the posts within that cell. We then summed the total number of posts in the entire dataset associated with any of those topics. For instance, if a cell included posts appearing in Topic 1 and Topic 2, we computed the cell's upper bound by summing all posts assigned to either topic. Since a topic may appear in multiple cells, this method intentionally overcounts, providing a conservative estimate of how widely a challenge or use case may occur. Figure 2 visualizes this mapping and is explained further in Section IV. To validate this estimation approach, we drew a new stratified random sample of 200 posts from the full dataset (2008-2024) and independently annotated them using the same methodology described in Section III-C. As can be seen in the extended supplementary material, both data distributions (625 original sample and 200 validation sample) are similar, suggesting that our initial topic-based upper bound estimates align with the proportions observed in a separate random sample.

E. Limitations

Despite our approach to developing and validating our keyword list for data extraction, we acknowledge the potential for false negatives. For example, while we removed keywords like "update" to avoid false positives, we may risk losing relevant posts about security-related updates in technical processes. We acknowledge that the LLM-based approach may miss out on further relevant keywords due to a bias in the responses provided by the AI model. We mitigated this risk in evaluating keywords against known security and privacy standards [90], [91]. Our dataset may not exhaustively cover all cloud-related S&P posts, but is comprehensive enough for a detailed analysis. We encourage future work in this space to validate our findings. Another limitation of our work is that our dataset captures the security and privacy architecture offered by the selected cloud providers up until 2024.

IV. Findings

Overall, we obtained 251,904 cloud related security and privacy posts ranging from 21 Aug. 2008 - 31 March 2024. 194,482 posts ($\approx$ 77.2%) have at least one answer, and 96,069 ($\approx$ 38.1%) posts have an accepted answer. Here, an *accepted answer* means one that the original posters have marked as the most helpful or appropriate solution to their problems [105]. Authentication, password, and security are the top three security and privacy keywords in our dataset, and amazon-web-services, azure, and firebase are the top three tags. Table I further shows the top ten tags, keywords, and their corresponding percentages. Figure 6 in Appendix D shows temporal plots for the number of questions asked and average answer counts over the years. Amazon Web Services, Microsoft Azure, Firebase, Google Cloud Platform, and Heroku are the top five most commonly discussed cloud providers (see Figure 3).

To support our findings, we provide links to questions and answers as PXXXX and AXXXX, respectively, where XXXX stands for the last four digits of the question/answer's ID (only the last four digits are presented for brevity).

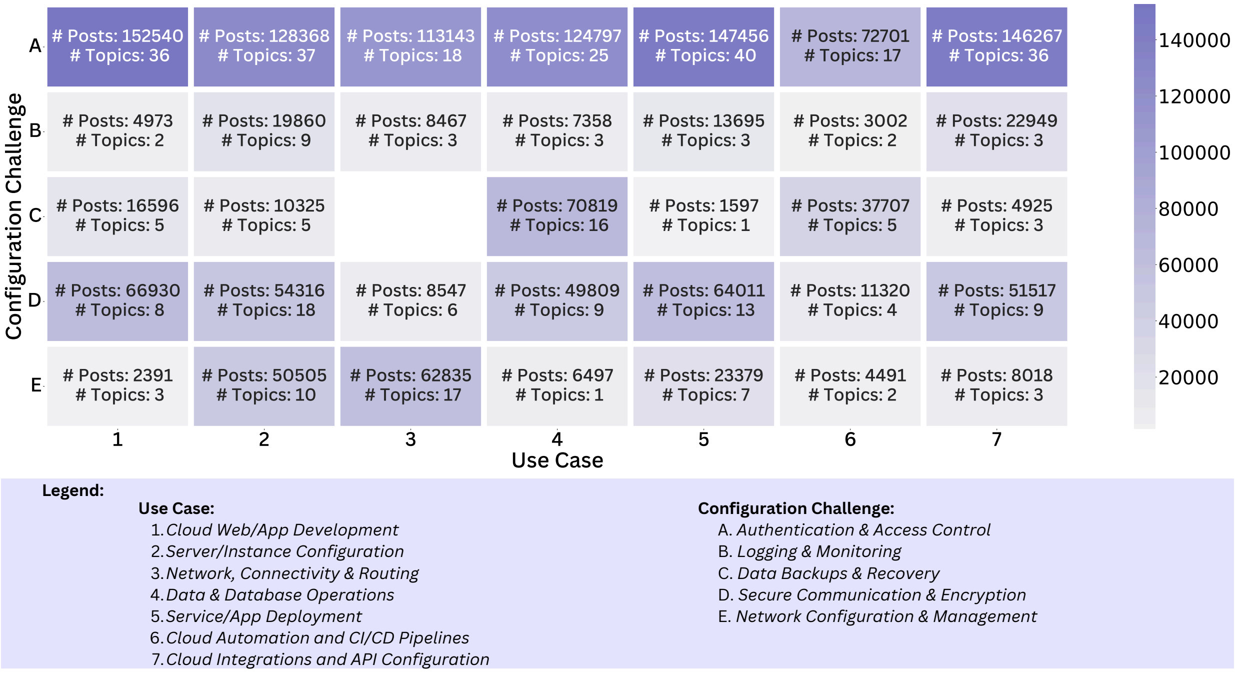

Configuration Challenge	1	2	3	4	5	6	7
A	# Posts: 152540 # Topics: 36	# Posts: 128368 # Topics: 37	# Posts: 113143 # Topics: 18	# Posts: 124797 # Topics: 25	# Posts: 147456 # Topics: 40	# Posts: 72701 # Topics: 17	# Posts: 146267 # Topics: 36
B	# Posts: 4973 # Topics: 2	# Posts: 19860 # Topics: 9	# Posts: 8467 # Topics: 3	# Posts: 7358 # Topics: 3	# Posts: 13695 # Topics: 3	# Posts: 3002 # Topics: 2	# Posts: 22949 # Topics: 3
C	# Posts: 16596 # Topics: 5	# Posts: 10325 # Topics: 5		# Posts: 70819 # Topics: 16	# Posts: 1597 # Topics: 1	# Posts: 37707 # Topics: 5	# Posts: 4925 # Topics: 3
D	# Posts: 66930 # Topics: 8	# Posts: 54316 # Topics: 18	# Posts: 8547 # Topics: 6	# Posts: 49809 # Topics: 9	# Posts: 64011 # Topics: 13	# Posts: 11320 # Topics: 4	# Posts: 51517 # Topics: 9
E	# Posts: 2391 # Topics: 3	# Posts: 50505 # Topics: 10	# Posts: 62835 # Topics: 17	# Posts: 6497 # Topics: 1	# Posts: 23379 # Topics: 7	# Posts: 4491 # Topics: 2	# Posts: 8018 # Topics: 3

Configuration Challenge
Use Case
140000
120000
100000
80000
60000
40000
20000
Legend:
Use Case:
1. Cloud Web/App Development
2. Server/Instance Configuration
3. Network, Connectivity & Routing
4. Data & Database Operations
5. Service/App Deployment
6. Cloud Automation and CI/CD Pipelines
7. Cloud Integrations and API Configuration
Configuration Challenge:
A. Authentication & Access Control
B. Logging & Monitoring
C. Data Backups & Recovery
D. Secure Communication & Encryption
E. Network Configuration & Management

Fig. 2: Mapping use cases to configuration challenges uncovered in the qualitative sample. This figure was obtained by mapping each post's assigned use case and configuration challenge. The count of posts in each cell is an upper bound value obtained via the methodology in Section III-D. Each cell also contains the total number of assigned topics of the respective posts in that cell. The extended supplementary material presents the actual topics per each cell.

Tag	Counts	SP Keyword	Counts
amazon-web-services	5.69%	authentication	15.80%
azure	4.18%	password	11.21%
firebase	4.02%	security	9.90%
firebase-authentication	2.36%	permissions	6.10%
google-cloud-platform	1.64%	https	5.21%
amazon-s3	1.60%	authorization	4.86%
amazon-ec2	1.47%	secret	4.79%
node.js	1.36%	iam	4.72%
javascript	1.31%	ssl	4.14%
android	1.26%	vpc	3.06%

TABLE I: Counts of the top 10 most frequent tags and S&P keywords in our total dataset, with counts as percentages of total tags and keywords respectively.

A. *Mapping Cloud Security & Privacy Challenges (RQ 1)*

We uncovered seven use cases of the cloud ecosystem and five recurring security and privacy configuration challenges between 2008 and 2024. To understand the broader prevalence of these challenges, we mapped the coded use cases and configuration challenges to the topics uncovered in our full dataset. While this offers an upper-bound estimate, since posts in a topic may not uniformly reflect the same challenges, it enables a high-level view of how these issues are distributed. Figure 2 presents this mapping. It highlights key intersections where specific use cases, e.g., application deployment or network setup, frequently co-occur with challenges like authentication or secure communication. These intersections reveal common pain points in cloud security and provide guidance for where support and documentation may be most needed. In this section, we present these use cases, configuration challenges, and how these map onto each other. We start by outlining the seven use cases:

1) **Cloud Web/App Development:** Posts discussed building web applications using cloud-based services like authentication APIs, credential managers, and databases. These posts focused on implementing app functionality e.g., user login flows, conditional UI redirection, and more.
2) **Server/Instance Configuration:** Posts discussed provisioning and managing virtual machines and server resources (e.g., storage, CPU, load balancers) on platforms like AWS EC2. Questions about environment setup, service migration, and lifecycle management of cloud resources were coded in this category.
3) **Network, Connectivity, and Routing:** Posts discussed setting up cloud networking components like firewalls, Virtual Private Clouds (VPCs), NATs, VPNs, DNS, and route tables. These posts addressed securing connections, managing IP-based access, and ensuring correct traffic routing across internal and external networks.
4) **Data and Database Operations:** Posts discussed deploying, configuring, and accessing cloud-hosted databases. These posts also asked questions related to security rule customizations, CRUD operations, and configuring backup, replication, and disaster recovery of data environments.

5) **Application and Service Deployment and Management:** Posts discussed deploying applications, APIs, or services to the cloud. Operators asked about handling secrets, managing SSL certificates, configuring runtime environments, and configuring post-deployment monitoring tools.
6) **Cloud Automation and CI/CD Pipelines:** Posts discussed automating deployment workflows and managing CI/CD configurations. Operators asked questions pertaining writing templates, debugging automation scripts, and reducing manual maintenance through infrastructure-as-code and DevOps practices.
7) **Cloud Integrations and API Configuration:** Posts discussed connecting disparate cloud tools and services using APIs. This included scheduling API calls, configuring notification/messaging systems, and rerouting API requests to improve reliability and performance.

Next, we elaborate on each security and privacy challenge category and describe how they manifest in specific use case contexts. For space reasons, we focus on prominent use cases (with more than 30,000 posts) and provide a description of the remaining ones in the supplementary material.

A. Authentication and Access Control: A prominent configuration challenge lies in managing authentication and fine-grained access control across cloud services. This category includes posts that describe issues involving authentication mechanisms, e.g., identity management, token scopes, and policy enforcement, that must be successfully configured prior to applying access controls. Notably, this excludes lower-level access mechanisms e.g., firewalls or security groups (compare Section IV-A). These challenges arise prominently across all seven identified use cases.

Use Case 1: In the context of web and app development, operators encounter hurdles when implementing end-user authentication flows using cloud identity services like Firebase Authentication or Azure Active Directory. Posts describe difficulties configuring login and signup mechanisms, integrating multi-factor authentication (MFA), or validating user input via email or SMS. E.g., P3757 discusses integrating Azure Active Directory Federation Services (ADFS) for enabling Single Sign-On (SSO) with Office 365, while P5660 highlights problems where Firebase Authentication fails to respond correctly to login attempts, regardless of password accuracy, prompting debugging of client-side logic.

Use Case 2: Beyond end-user flows, access control misconfigurations surface in use cases involving instance configuration and cloud integrations. These issues relate to configuring IAM policies, token scopes, or environment-specific credentials that enact privacy. E.g., P7010 and P1866 involve configuring programmatic access to backend services or AWS accounts using scoped tokens and SSO scripts.

In these posts, incorrect permission setups frequently lead to authorization errors, e.g., HTTP 401 (unauthorized), 403 (forbidden), or Cross-Origin Resource Sharing (CORS)-related failures. These are reported in various environments, including VPCs, databases, and deployment pipelines. For example, P5073 describes a case where an operator cannot access AWS S3 buckets due to insufficient IAM privileges, leading to confusion over policy rules.

Use Case 3: Authentication and access control issues appear during network configuration and connectivity. E.g., developers face difficulties enabling RDP access to cloud subnets, managing IP allowlists, or troubleshooting connectivity within VPCs. Common challenges include: inability to enable remote desktop access to cloud subnets (e.g., Azure), failures when connecting to cloud-hosted instances via external IPs, managing geolocation-based traffic control in distributed deployments, errors in configuring VPC-level access rules or DNS-based redirection, and connectivity timeouts or permission denials within VPCs due to misconfigured credentials.

Use Case 4: In the context of data management, operators report permission issues when connecting to cloud-hosted databases, especially in multi-service architectures where one cloud component must authenticate to access another. These challenges are typically rooted in identity misconfiguration or improperly scoped tokens used to authenticate database queries or automated tasks.

Use Case 5: During application deployment, permissions-related issues arise post-deployment, particularly when services cannot access required resources e.g., environment variables, secrets, or storage systems. E.g., operators report inaccessibility of cloud services due to missing runtime permissions, and errors during server-side rendering caused by improperly scoped credentials or token mismanagement.

Use Case 6: Within DevOps and cloud automation, authentication and permission-related challenges arise in the form of access-denied errors when automation and CI/CD jobs need to access resources or connect across services. E.g., P3068 details a case where an operator seeks to grant CI/CD services access to a code repository stored on a cloud server.

Use Case 7: Within API configuration and cloud integrations, operators encounter authentication barriers while connecting services across cloud platforms. These include token misconfigurations, CORS violations, or unexpected request failures when invoking secured endpoints. For example, operators report: API request failures due to invalid scopes or CORS misconfiguration, permission-related service shutdowns triggered by incorrect token usage, errors establishing secure service-to-service connections across platforms, challenges integrating cloud-native APIs with hosted front-end or middleware components, and more.

> **Key Takeaway**
>
> Authentication and access control serve as a foundational layer for all cloud configurations. Misconfigurations at this layer can cascade into failures across other domains e.g., networking, deployment, and development, reinforcing the need for a holistic understanding of identity and access mechanisms in secure cloud systems.

B. Logging and Monitoring: This challenge focuses on configuring and maintaining logging and monitoring services in cloud environments to support security, error detection, and performance analysis. Operators struggle with ensuring that logging services configured in one environment (e.g., a third-party monitoring system or external API) correctly interface with cloud-native platforms, facing issues e.g., data format inconsistencies. Posts describe persistent difficulties with log extraction, transmission, and visualization when using tools e.g., AWS CloudTrail, AWS CloudWatch, Azure Log Analytics, and Google Cloud Operations. Operators commonly seek guidance on customizing logging pipelines to meet specific operational requirements. E.g., in P9241, an operator attempts to create a rule within AWS CloudWatch to identify terminated instances by name so they can be removed from an external monitoring system. Similarly, P0149 involves an operator trying to add timestamps to logs retrieved from CloudWatch, while P8219 centers on integrating security logs using the Azure Log Analytics visual interface. Networking-related posts focus on the integration of logging frameworks across services connected within Virtual Private Clouds. In the context of application deployment, operators encounter performance bottlenecks or failures when synchronizing system events with dashboards or 3rd-party monitoring tools.

Key Takeaway

A recurring concern in logging and monitoring is integrating logging in different platforms and consolidating logs from diverse sources, e.g., serverless functions, virtual machines, and container orchestration platforms, only to encounter inconsistent log formats, missing metadata, or configuration incompatibilities.

Category	Description and Examples	Count
Official Documentation	Official documentation pages. Examples: docs.microsoft.com and docs.aws.amazon.com	16,349 (24.92%)
Code Sample	Official code repositories on Github: github.com	13,096 (19.97%)
Hosting Provider Website	Websites of cloud providers. Examples: firebase.google.com and aws.amazon.com	12,260 (18.69%)
Discussion Forum	Unofficial discussion forums. Examples: stackoverflow.com and serverfault.com	11,209 (17.09%)

TABLE II: The top four domain categories of the top 63 domains (count $\geq$ 100) used in our entire dataset.

C. Data Backups and Recovery: Questions in this category focus on the configuration, automation, and security of backup and disaster recovery mechanisms in cloud environments.

Use Case 4: Posts frequently express concern about secure data restoration. Operators often seek guidance on how to replicate critical infrastructure across regions, schedule database snapshots, or maintain data availability under failure scenarios. They ask which parts of the backup workflow require encryption, how to restrict restoration privileges to authorized personnel, and how to avoid unintentional data exposure during the recovery process. Such questions underline the intersection of resilience planning and data security, as the failure to adequately secure or test restoration paths may transform a backup into a vulnerability rather than a safeguard.

Use Case 6: In cloud automation contexts, operators encounter issues with automating backup processes for virtual machines or cloud-hosted databases, particularly when using infrastructure-as-code tools like AWS CloudFormation. Problems also arise when synchronizing data across hosting providers or managing update consistency during failover transitions. These challenges reflect the nuanced task of maintaining redundancy and state fidelity across environments.

Key Takeaway

Configuring backups and recovery procedures is not a purely operational task, but one entangled with automation, cloud infrastructure design, and access control. Ensuring recoverability while minimizing downtime, cost, and risk demands expertise across all these domains.

D. Secure Communication and Encryption: This challenge category encompasses posts involving the configuration of encryption mechanisms to enact privacy and secure data transfer practices within cloud environments. The posts illustrate ongoing difficulties in ensuring encrypted data transmission, managing certificates, and maintaining end-to-end security compliance across cloud services.

Use Case 1: Within web and app development, a recurring challenge involves the transition from HTTP to HTTPS for securing browser-server communication. Posts frequently address the addition and renewal of SSL certificates for web applications hosted on cloud infrastructure. In particular, operators report configuration difficulties related to HTTPS enablement using cloud-native services e.g., AWS Load Balancers. E.g., P9723 describes an attempt to configure HTTPS access for a website hosted on AWS via a Load Balancer, where the operator seeks guidance on ensuring secure access and certificate setup. Additional challenges in this context include SSL handshake failures and certificate validation errors.

Use Case 2: Encryption-related issues in server and instance configuration commonly stem from deploying secure communication protocols for internal and external service interactions. Posts detail challenges e.g., configuring HAProxy for API Gateway client certificate validation or integrating certificate-based authentication in self-hosted environments. Operators discuss the difficulties of implementing secure configuration practices when migrating from managed to self-hosted cloud deployments, where secret storage and certificate injection mechanisms are not abstracted by the provider. These misconfigurations can result in degraded service reliability and insecure data exposure.

Use Case 4: Operators discuss encrypting databases and files. Challenges arise in the shape of errors while integrating encryption utilizing cloud secret managers, file corruption, library limitations when handling large data, and more. For example, P4400 asks for help in encrypting specific database columns in Azure SQL Server using Azure Key Vault, asking whether column-specific encryption is feasible or not.

Use Case 5: During application deployment, posts reveal critical challenges in securing runtime environments through SSL configuration and secret management. Operators struggle to automate the insertion of SSL certificates into cloud-hosted applications, and report difficulties managing certificate expiration, CORS headers, and encryption libraries across platforms. Moreover, the handling of application secrets during deployment proves error-prone; failures in storing, formatting, or injecting secrets, e.g., database connection strings or API credentials, often lead to runtime failures or insecure deployments. These issues appear frequently in serverless functions and containerized applications, where manual certificate management may be insufficient or misaligned with deployment workflows.

Use Case 7: Cross-platform API integrations introduce additional encryption challenges, particularly when enabling SSL-secured communications between microservices or external clients. Several posts involve misconfigurations when securing service access through API Gateways using HTTP proxies or implementing SSL encryption between backend services. E.g., post P8087 details an SSL handshake failure when a middleware attempts to access a backend cloud-hosted API due to an untrusted certificate, prompting a request for guidance on middleware configuration. These scenarios illustrate the critical role of trust chains, certificate authorities, and TLS negotiation in achieving secure API interoperability.

Key Takeaway

Secure communication and encryption challenges highlight the complex interplay between certificates, secrets, and encryption libraries across all layers of cloud applications. Whether encrypting data in transit between services or securing data at rest in cloud databases, operators face difficulties in selecting appropriate tools, configuring them correctly, and maintaining their security posture over time.

E. Network Configuration and Management: This challenge category centers on configuring and managing networking components within cloud environments. Unlike authentication-based challenges, these posts primarily involve network-layer concerns where access control does not depend on user identity verification, but rather on IP filtering, routing, or infrastructure topology.

Use Case 2: In the context of server and instance configuration, posts reveal common difficulties in achieving secure and reliable connectivity across distributed environments. Operators report errors when launching instances in different availability zones or configuring Terraform scripts for creating security groups. Other questions focus on network reachability problems between on-premise servers and cloud-hosted virtual machines. These configurations often require correctly defined network security groups and routing logic to ensure desired data flows. Missteps in these areas can manifest as failed pings, inaccessible services, or unintended exposure of private endpoints.

Use Case 3: Posts in this use case focus on core networking components e.g., VPC peering, NAT gateways, route tables, firewalls, VPN tunnels, and DNS services like AWS Route 53. Operators frequently report failures in setting up site-to-site VPNs, configuring DNS failover mechanisms using services like AWS Route 53, or managing subnet isolation and IP assignments. Firewalls, in particular, appear often as tools to limit traffic based on IP origin or port—without involving identity-based authentication—but pose challenges when used in dynamic environments. A recurring theme is the burden of repeatedly updating firewall or access control rules in response to changing infrastructure, e.g., ephemeral IPs associated with load balancers or serverless services.

Misconfigurations in these networking elements can easily disrupt inter-service communication, as even small errors in a route table or NAT configuration can block legitimate traffic or expose sensitive endpoints. E.g., post P5969 describes an attempt to configure a firewall that restricts access to a self-hosted web service, allowing only a cloud-hosted website to connect. However, the website's IP address changes frequently due to its placement behind a load balancer, leading the operator to seek solutions that avoid manual rule updates while maintaining security guarantees. These challenges highlight the complexity of operating secure, reliable cloud networks in dynamic and distributed environments.

Key Takeaway

Network configuration and management challenges emerge as foundational concerns that intersect with deployment, data access, and service integration. Ensuring that traffic flows securely and predictably across regions and services remains a non-trivial task, particularly in multi-tenant and multi-cloud architectures.

B. Human-Centric Issues in Cloud S&P (RQ 2)

Our qualitative analysis uncovered additional challenges and reasons why people ask questions on Stack Overflow when it comes to configuring cloud security and privacy. These issues help contextualize and identify potential reasons why people struggle in implementing and configuring cloud use cases, which lead to errors and misconfiguration.

1) Challenges in Using Help Resources: This section presents an overview of the usage of help resources from 2008-2024. We did not observe any major temporal trends. In total, 92,129 posts (approximately 35.6% of the dataset) referenced external domains related to cloud S&P. The top 63 domains, which appear more than 100 times, account for 61,767 posts (approximately 24.5%). These posts often contain multiple domain references, with 21,637 posts (8.59%) citing more than one domain. The most frequently cited domains include: *github.com*, *stackoverflow.com*, *docs.aws.amazon.com*. The most common domains are summarized in Figure 7. Table II categorizes the 63 most frequently referenced domains, which together cover about 81.1% of the dataset. For a detailed listing of all 63 domains see the extended supplementary material.

Qualitative analysis helped us understand the context behind using the help resources. A common theme is that existing help resources often fail to provide tailored guidance for specific use cases. Operators express frustration with documentation examples that either do not work, are not aligned with their requirements, or need to be expanded to cover specific configurations. E.g., P7745 encounters authorization errors when attempting to follow official documentation to connect to Dropbox through a web app hosted in the cloud. Similarly, in P6529, an operator is unable to find clear instructions in the Firebase documentation on linking multiple phone numbers to the same account, and no answers are available.

Other posts reveal confusion trying to understand the content of help resources. In some cases, they seek clarification on whether provided code examples will work in their specific environment or with their custom configurations. An example is P6288, where the operator is unable to understand the API authentication process in the context of integrating Laravel with AWS. Additionally, questions frequently ask for links to relevant documentation that would help them understand the relationships between components in complex cloud services, e.g., AWS EC2 instances (P4704). Many of these issues are also seen where operators disclose they are new to the platform or cloud service they are working with.

2) Unusable Services: We uncovered these sub-themes:

Unusable Code Libraries and Commands: Questions involve errors when using cloud-related code libraries or command-line tools. These issues often arise when operators attempt to copy-paste or adapt code examples from help resources without fully understanding the underlying implementation details. A primary example of unusable code libraries are cloud authentication libraries, e.g., Firebase Authentication, that have numerous questions related to implementing authentication for software applications. Primarily, questions ask for help in debugging errors in code that cause unresponsiveness or incorrect outputs from the library. An example is P3363, where the operator is getting malformed credentials while trying to programmatically integrate Facebook authentication in a Unity app using Firebase. There are questions on how to use these libraries to implement specific use cases, e.g., SMS and email verification, or to debug errors, e.g., to see if authentication is initialized correctly or not. People ask for help in understanding and accessing the various callbacks and parameters provided by these libraries. There are mentions of issues in accessing and utilizing authentication tokens. Other examples discuss configuring servers using templates or writing access control policies and security rules. Various posts face issues when using code for data/ file retrieval/ storage from cloud storage.

Unusable Visual Interfaces of Cloud Services: Another challenge is the complexity and usability issues with the visual interfaces provided by cloud services. These interfaces, e.g., dashboards and settings menus, are often used to configure crucial components like logging mechanisms, permissions, authentication tokens, and security groups. However, questions frequently encounter difficulties when attempting to select the correct options within these interfaces. In P3718, the operator experiences errors while configuring a Virtual Private Cloud (VPC) to allow traffic to a MongoDB database using the Azure interface. Similarly, P4418 seeks guidance on implementing OWASP recommendations within the Azure project's UI.

3) Knowledge Gaps: In addition to usability challenges, many struggle with knowledge gaps that hinder their ability to make informed decisions or fully understand the systems they are working with. These gaps often manifest as questions seeking conceptual knowledge or specific implementation details about cloud services:

Understanding How Systems Work: Operators seek guidance on understanding how different cloud services function and interact. This knowledge is essential for configuring cloud systems correctly and troubleshooting potential issues. E.g., P9872 asks how AWS ensures Elastic Block Store isolation so that only the project owner can see or access the block storage instance.

Secure Coding/ Configuration Practices: Operators ask questions related to secure coding practices to protect cloud implementations from potential security threats. E.g., P6399 is concerned about specific AWS's recommendations regarding DNS settings as the operator fears that they may expose their domain to unauthorized use. These questions may ask for conceptual and implementation details to ensure security. One example is P4787, where the operator asks whether their written client-side code logic to query files from their Firestore database is sufficient to protect file access or if a malicious user could exploit it to gain unauthorized access.

Testing-related Queries: Questions focus on how to test applications, cloud services, or APIs to ensure they meet security and performance expectations. These queries often reveal uncertainty about the most appropriate testing strategies for cloud-hosted systems. For e.g., P4151, where the operator evaluates Azure VMs for their organization's disaster recovery.

Key Takeaway

Despite widespread referencing of official documentation and external help resources, operators frequently encounter substantial challenges in applying this information to their specific cloud S&P use cases. These challenges stem from documentation that lacks contextual guidance, the complexity and limited usability of both code libraries and visual interfaces, and critical knowledge gaps in secure configuration practices. Together, these issues reveal a persistent disconnect between available support resources and the practical realities of cloud development, underscoring the need for more actionable, comprehensible, and context-aware support mechanisms.

C. Accepted Answers (RQ 3)

The accepted answers offer tailored advice to help resolve the questions by adopting three main strategies to resolve or work around the bottlenecks operators face:

1) Code Changes and Parameter Explanation: The first strategy involves recommending code changes, e.g., adding or editing code attributes, configuring environment parameters, or guiding on how to use specific APIs. E.g., one answer (A9367) provides a code example that helps an operator migrate from Google Sign-In to Firebase Authentication, addressing common implementation challenges in this context. This strategy is often employed when operators encounter errors due to incorrect or outdated code snippets, and answers aim to clarify how to modify or replace these elements.

2) Cloud Service Usage: The second strategy suggests alternative cloud services or tools to address the question's issue. In these cases, answers typically begin by explaining the problem and then recommend a more efficient or effective approach using different cloud operations or services. E.g., P5783 asks how to retrieve the deployment password for an Azure website. The accepted answer (A8894) provides instructions on how to obtain the required credentials from the Azure Portal, including relevant screenshots of the dashboard to guide the operator through the process.

3) Reference to Official Documentation and Help Resources: In addition to code modifications and service recommendations, many answers also refer to official documentation or external help resources. This approach provides more detailed reading material and examples to resolve their questions independently. E.g., A2114 points to official AWS documentation on extracting timestamps from their logging and monitoring service, helping the operator better understand and implement the required functionality.

4) Meta Analysis of Help Resources: A meta-analysis of the domains referenced in accepted answers shows that these resources are consistently drawn from a small set of trusted sources. Across the larger dataset, a total of 44,879 accepted answers (approximately 46.7%) refer to specific domains. The top 48 domains, which appear more than 100 times, account for the majority of these references. The most frequently cited domains include: *docs.aws.amazon.com*, *github.com*, *docs.microsoft.com*, *firebase.google.com*, *cloud.google.com*, *stackoverflow.com*. These domains largely overlap with those found in the original questions, suggesting a consistent reliance on well-established documentation sources.

Key Takeaway

Accepted answers primarily address configuration challenges through three complementary strategies: code changes and explanations, service usage suggestions, and references to official documentation. The consistent reliance on a set of trusted help resources, often mirrored in the questions themselves, highlights a shared ecosystem of knowledge that operators and answerers depend on to resolve cloud configuration problems.

V. Discussion

While prior research has focused on usability challenges in security configuration [53], [54], [55], [56], operator misunderstanding [50], [51], [16], [52], [17] and characteristics [18], or documentation use [57], these papers examine specific user groups or security aspects in siloed contexts. In contrast, our study illustrates how security and privacy challenges routinely surface across the entire spectrum of cloud development, regardless of the underlying technology. From configuring authentication for a web application to debugging broken CI/CD deployments or managing third-party service integrations, operators face cross-cutting, persistent responsibilities.

Our dataset includes both narrowly scoped debugging questions and broad inquiries about the feasibility of certain cloud configurations—reflecting how the cloud landscape fosters not only technical missteps but also fundamental uncertainty. The prevalence of such questions underscores that being stuck is not the exception but the norm, and current tools are failing to meet operator needs and expectations. We argue for three complementary paths forward: (i) addressing the usability of cloud security and privacy systems, (ii) expanding and adapting help and training resources to reflect the evolving and context-dependent nature of operators' challenges, and (iii) studying the impacts of organizational security culture.

A. Usability Challenges in Cloud Security and Privacy

Our first recommendation is to improve the usability of core cloud security and privacy mechanisms to mitigate configuration pitfalls at their source. A starting point in this direction, is the use case of authentication and access control configuration. Our analysis highlights authentication and access control as a particularly fraught area—appearing across all seven use cases in our study and emerging as the most common challenge in our analysis (Figure 2). Authentication and access control is foundational to cloud security, yet its configuration remains difficult and error-prone. This prevalence makes authentication and access control a unique case: whereas other security and privacy discussions on Stack Overflow arise from secondary concerns (e.g., deployment friction or API connectivity), it is both a primary goal and a persistent pain point. Operators must configure login flows, token scopes, policy permissions, and cross-service identity bindings, all of which demand a nuanced understanding of security concepts and service-specific APIs.

To ground the above observations about authentication and access control, we use one representative case to show how context-insufficient guidance can turn "correct" examples into deployment-time roadblocks. In P5073 (Section IV-A), an operator relied on a provider example (e.g., IAM policy examples) for S3 bucket access but still could not access the bucket due to insufficient privileges, specifically, confusion about which execution role the policy should be attached to and how policy scope is evaluated. In that Stack Overflow thread, the accepted answer clarifies these contextual assumptions rather than changing the JSON snippet itself (i.e., the example was syntactically correct, but attached to the wrong role and scope). This pattern generalizes: operators need (i) *role-aware starter templates* that distinguish CI/CD identities from runtime services; (ii) *decision trees* that guide where to attach which permissions and under what conditions; and (iii) *in-context checkers* that flag likely mis-scoping (e.g., a policy attached to a role that never assumes the relevant principal or lacks required actions/conditions).

Building on this pattern, despite longstanding work on improving authentication usability (e.g. [106], [107], [108], [109], [110]), we recommend future research to revisit the usability of authentication tooling from the operator's view where integrating, debugging, and securing authentication happens in the context of larger, interdependent systems. We recommend more user-centric evaluations of these tools in practice, focusing on integration workflows, error recovery, and failure visibility across the lifecycle of cloud systems. One promising starting point is Firebase Authentication as this was the most prominently discussed library in our dataset.

B. Help Resources and Education

Improving cloud service usability is necessary but insufficient on its own. As our findings show, cloud operators encounter a wide range of evolving and context-specific challenges that even the best-designed services cannot always anticipate. We observed a spectrum of questions: some narrowly focused on error messages (e.g., P3350), while others asked whether particular architectures or configurations were even feasible (e.g., P3408). This diversity shows that help-seeking is not just simply about *how to fix something*, but also about *how to reason through* what is possible in a constantly shifting cloud landscape. Thus, our second recommendation is to re-imagine help-seeking as a core part of secure cloud operation. This step is important as prior work has uncovered that these help resources impact developers' security and privacy decisions [19], [20]. Further, we discuss the role of certifications as a form of security education.

1) Documentation: Prior work has studied how to make documentation more usable (e.g. [57], [58], [111]). Recent propositions, e.g., making documentation interactive [58], are promising but still require further research into how such designs would play out in the wild. Our study suggests that existing documentation often lacks specific implementation details tailored to the individual use cases discussed on Stack Overflow questions (IV-B1). Future work should study how documentation can be made more context-specific to meet tailored requirements. A promising direction would be to tailor documentation based on specific documentation usage patterns, that prior work has identified depends on users' background, context, and prior experience with the service [57].

2) LLMs: People have recently shifted to seeking help from LLM chatbots, e.g., ChatGPT, for technical guidance. This reliance on AI models has also resulted in a decline in the usage of Stack Overflow, as also evident by a smaller number of posts in our dataset in 2023-24. LLMs are a promising opportunity for providing support in cloud configurations. However, the effectiveness and accuracy of answers from these AI models are still questionable, as indicated by prior research: people find human-generated answers to be more correct, concise, and useful than LLM-generated content [71]. This issue may be because these models lack the context behind the questions users have: users may unknowingly not provide necessary information. For example, in P1434, the question asks on automating incremental backups of Amazon server instances but does not give more context on specific implementation details of their cloud setup and the nature of these required incremental backups, which may be imperative to answer the question. Obtaining accurate answers from ChatGPT not only relies on building better models but also on effective prompt engineering from users which may be lacking. Future work should investigate how fine-tuned LLMs can be integrated and deployed within operators' cloud configurations to provide context-sensitive help when required. These LLMs would be fully aware of the capabilities and limitations of operators' setups and hence offer more contextualized advice. Future work should also focus on curating accurate training data for these LLMs to avoid training on the potentially erroneous data.

3) Cloud Security Certifications: A brief review of two major certification programs [112], [113] reveals that their coverage may not fully reflect real-world operational experiences. We could not fully explore one of the certifications as it required paid subscriptions for certain modules, while the second certification consisted primarily of textual material and exam-based assessments, rather than scenario-driven, hands-on training. These certifications may not adequately prepare operators in real-world security configuration tasks which may have additional time and infrastructural constraints. Prior studies on system administrators further supports this observation: administrators do not typically attribute their expertise to formal education but rather to experiential, "hands-on" learning [61], and they often feel insufficiently trained for real-world system administration [114]. It is imperative for future work to conduct an in-depth examination of the effectiveness of certification programs, as it can inform the design of more effective training curricula that bridge the gap between formal education and the realities of daily cloud security administration.

C. Organizational Security Culture

One limitation of this work is that there is no reliable way to link question-askers on Stack Overflow to the organizations they work in, as they do not mention their organizations systematically or discuss their organizational culture and workload in posts. However, the impact of organizational culture on security outcomes is a promising direction for future research. We discuss two angles relevant to operators below:

1) Internal Knowledge Bases: Prior work in the domain of system administration, especially in system updates and patch management, has explored how cloud operators rely on internal knowledge bases, e.g., organizational policies, formal update processes, and team members to seek help and exchange information [61], [115], [51], [114]. E.g., prior work has shown that internal knowledge bases are preferred over official manufacturer support channels for troubleshooting and patch updates [61], [51], [114]. However, the effectiveness of these internal knowledge bases may be impaired due to management constraints, lacking support infrastructure in smaller teams, and other factors [51]. How operators approach internal knowledge bases in the context of security and privacy, along with the exact nature of the challenges they face in seeking help from internal knowledge bases, is understudied. An important direction for future work is to understand how cloud operators collectively build, maintain, and rely on a mix of information sources, as this would reveal how organizational culture and informal learning practices shape security-related decisions and troubleshooting in cloud operations.

Future work should also conduct longitudinal studies with cloud operators to further explore how findings of our work correlate with the security maturity and posture of organizations. Specifically, it would be valuable to investigate how changes in provider APIs, tooling, and organizational culture influence the emerging challenges faced by cloud operators over time.

2) Social Embeddedness: The organizational context in which these cloud configuration technologies are used is also essential. Operators working in large-scale organizations may have more avenues to seek technical help, e.g., from other developers, compared to self-hosters, who often engage in social constellations to maintain IT infrastructure [17]. An interesting example is from May 2024, when UniSuper's cloud operators collectively worked with Google Cloud to restore UniSuper's accidentally deleted Google Cloud setup configuration and customer data [116]. Understanding the social context in which cloud technologies are used is important to leveraging these organizational and social factors to provide streamlined help. Future work should investigate how security advice for cloud operators can be disseminated by leveraging the organizational context in which cloud services are used.

VI. Conclusion & Future Work

Cloud security and privacy misconfigurations are not isolated edge cases—they are recurring challenges that span the entire cloud ecosystem. Our mixed-methods analysis of over 250K Stack Overflow posts uncovered seven common use cases and five cross-cutting security and privacy challenges, highlighting the persistent complexity operators face when deploying, integrating, and managing cloud services. These challenges stem not only from technical errors but from broader usability failures and insufficiently tailored support systems. Future studies should focus on understanding the real-word configuration experiences of operators in each of the use-cases identified by our study, their help-seeking mechanisms, and their experiences of reporting unhelpful official documentation on cloud platforms to gain a deeper understanding of this space. Addressing the configuration challenges requires improving the design of security and privacy tools and feedback mechanisms—especially around authentication and access control—and expanding context-aware, socially and organizationally informed help resources. Future research is needed on designing and evaluating LLM solutions grounded in results from operator experiences and usability studies.

Ethics Considerations

There are no anticipated ethical concerns. Our institutions do not require IRB review for studies using publicly available online data. We use Google BigQuery's and Stack Overflow Data Dump's publicly available datasets. In line with community norms such as the ACM Code of Ethics [117] and the Menlo Report [118], the SQL queries we wrote to extract data did not collect or analyze fields that may uniquely identify users, such as their usernames or user IDs. However, we acknowledge that users may deanonymize themselves in the posts. While we did not encounter such a case, we ensured not to code or report on any personal information.

Acknowledgments

We thank the reviewers for their invaluable feedback, which helped us improve our paper. This work was carried out under CIRCLE: The CISPA x LUMS Research Partnership on Internet For Everyone. Sumair Ijaz and Shafay Kashif were supported by the German Academic Exchange Service (DAAD) grant number 57708564 while working as undergraduate research assistants on this project at LUMS. Lea Gröber conducted a portion of this work while supported by a fellowship within the IFI programme of DAAD; funded by the Federal Ministry of Education and Research (BMBF).

References

[1] Amazon Web Services, "Aws website," https://aws.amazon.com/, accessed: 2024-01-26.

[2] Google Cloud Platform, "Gcp website," https://cloud.google.com/?hl=en, accessed: 2024-01-26.

[3] Microsoft Azure, "Azure website," https://azure.microsoft.com/en-us/, accessed: 2024-01-26.

[4] I. A. T. Hashem, I. Yaqoob, N. B. Anuar, S. Mokhtar, A. Gani, and S. U. Khan, "The rise of "big data" on cloud computing: Review and open research issues," *Information systems*, vol. 47, pp. 98–115, 2015.

[5] Flexera, "State of the cloud report 2024," https://resources.flexera.com/web/pdf/Flexera-State-of-the-Cloud-Report-2024.pdf, accessed: 2025-13-03.

[6] Thales Group, "2024 thales cloud security study - global edition," 2024, accessed:2024-31-01. [Online]. Available: https://cpl.thalesgroup.com/sites/default/files/content/CLOUD_AMI_pages/2023/2023-cloud-security-study-global-edition.pdf

[7] ——, "2024 thales cloud security study - global edition," 2024, accessed: 2025-03-13. [Online]. Available: https://cpl.thalesgroup.com/sites/default/files/content/cloud-security/2024/2024-thales-cloud-security-study-global-edition.pdf

[8] Trend Micro, "The most common cloud misconfigurations that could lead to security breaches," https://www.trendmicro.com/vinfo/us/security/news/virtualization-and-cloud/the-most-common-cloud-misconfigurations-that\-could-lead-to-security-breaches, accessed: 2024-31-01.

[9] Z. Whittaker. (2024) At&t says criminals stole phone records of 'nearly all' customers in new data breach. Accessed: 2024-07-18. [Online]. Available: https://techcrunch.com/2024/07/12/att-phone-records-stolen-data-breach/

[10] C. Apurva Venkat, "Cloud misconfiguration causes massive data breach at toyota motor," https://www.csoonline.com/article/575483/cloud-misconfiguration-causes-massive-data\-breach-at-toyota-motor.html,
Accessed: 2024-31-01.

[11] D. Leussink and R. Kantaro Komiya, "More than 2 million toyota users face risk of vehicle data leak in japan," https://www.reuters.com/business/autos-transportation/toyota-flags-possible-leak-more-than-2-mln-\users-vehicle-data-japan-2023-05-12/,
Accessed: 2024-31-01.

[12] CSO, "Cloud assets have 115 vulnerabilities on average — some several years old," https://www.csoonline.com/article/4003365/cloud-assets-have-115-vulnerabilities-on-average-some-several-years-old.html, accessed: 2024-01-26.

[13] Orca Research Pod, "2025 state of cloud security report," Aug. 2025, accessed: 2025-08-04. [Online]. Available: https://orca.security/lp/2025-state-of-cloud-security-report/#the-report

[14] O. Alrawi, C. Zuo, R. Duan, R. P. Kasturi, Z. Lin, and B. Saltaformaggio, "The betrayal at cloud city: An empirical analysis of {Cloud-Based} mobile backends," in *28th USENIX Security Symposium (USENIX Security 19)*, 2019, pp. 551–566.

[15] Z. Xiao and Y. Xiao, “Security and privacy in cloud computing,” *IEEE communications surveys & tutorials*, vol. 15, no. 2, pp. 843–859, 2012.

[16] C. Dietrich, K. Krombholz, K. Borgolte, and T. Fiebig, “Investigating system operators’ perspective on security misconfigurations,” in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, 2018, pp. 1272–1289.

[17] L. Gröber, R. Mrowczynski, N. Vijay, D. A. Muller, A. Dabrowski, and K. Krombholz, “To cloud or not to cloud: A qualitative study on self-hosters’ motivation, operation, and security mindset,” in *32nd USENIX Security Symposium (USENIX Security 23)*, 2023, pp. 2491–2508.

[18] L. Gröber, S. Lenau, R. Weil, E. Groben, M. Schilling, and K. Krombholz, “Towards privacy and security in private clouds: A representative survey on the prevalence of private hosting and administrator characteristics,” in *33rd USENIX Security Symposium (USENIX Security 24)*, 2024, pp. 6057–6074.

[19] Y. Acar, M. Backes, S. Fahl, D. Kim, M. L. Mazurek, and C. Stransky, “You get where you’re looking for: The impact of information sources on code security,” in *2016 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2016, pp. 289–305.

[20] F. Fischer, K. Böttinger, H. Xiao, C. Stransky, Y. Acar, M. Backes, and S. Fahl, “Stack overflow considered harmful? the impact of copy&paste on android application security,” in *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2017, pp. 121–136.

[21] A. S. Exchange, “Was chatgpt trained on stack overflow data?” https://ai.stackexchange.com/questions/38660/was-chatgpt-trained-on-stack-overflow-data, accessed: 2024-01-26.

[22] OpenAI, “Chatgpt,” https://chat.openai.com/, accessed: 2024-01-26.

[23] M. Coutinho, L. Marques, A. Santos, M. Dahia, C. França, and R. de Souza Santos, “The role of generative ai in software development productivity: A pilot case study,” in *Proceedings of the 1st ACM International Conference on AI-Powered Software*. New York, NY, USA: Association for Computing Machinery, 2024, p. 131–138.

[24] D. Molnar and S. E. Schechter, “Self hosting vs. cloud hosting: Accounting for the security impact of hosting in the cloud.” in *WEIS*, vol. 2010, 2010, pp. 1–18.

[25] W. Hu, T. Yang, and J. N. Matthews, “The good, the bad and the ugly of consumer cloud storage,” *ACM SIGOPS Operating Systems Review*, vol. 44, no. 3, pp. 110–115, 2010.

[26] D. Svantesson and R. Clarke, “Privacy and consumer risks in cloud computing,” *Computer law & security review*, vol. 26, no. 4, pp. 391–397, 2010.

[27] N. Syynimaa and T. Viitanen, “Is my office 365 gdpr compliant?: Security issues in authentication and administration,” in *International Conference on Enterprise Information Systems*. SCITEPRESS Science And Technology Publications, 2018.

[28] S. Subashini and V. Kavitha, “A survey on security issues in service delivery models of cloud computing,” *Journal of network and computer applications*, vol. 34, no. 1, pp. 1–11, 2011.

[29] M. Almorsy, J. Grundy, and I. Müller, “An analysis of the cloud computing security problem,” *arXiv preprint arXiv:1609.01107*, 2016.

[30] S. Iqbal, M. L. M. Kiah, B. Dhaghighi, M. Hussain, S. Khan, M. K. Khan, and K.-K. R. Choo, “On cloud security attacks: A taxonomy and intrusion detection and prevention as a service,” *Journal of Network and Computer Applications*, vol. 74, pp. 98–120, 2016.

[31] N. V. Juliadotter and K.-K. R. Choo, “Cloud attack and risk assessment taxonomy,” *IEEE Cloud Computing*, vol. 2, no. 1, pp. 14–20, 2015.

[32] A. Singh and K. Chatterjee, “Cloud security issues and challenges: A survey,” *Journal of Network and Computer Applications*, vol. 79, pp. 88–115, 2017.

[33] X. Wang, Y. Sun, S. Nanda, and X. Wang, “Credit karma: Understanding security implications of exposed cloud services through automated capability inference,” in *32nd USENIX Security Symposium (USENIX Security 23)*, 2023, pp. 6007–6024.

[34] A. Mehta, M. Alzayat, R. De Viti, B. B. Brandenburg, P. Druschel, and D. Garg, “Pacer: Comprehensive network {Side-Channel} mitigation in the cloud,” in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 2819–2838.

[35] S.-W. Li, J. S. Koh, and J. Nieh, “Protecting cloud virtual machines from hypervisor and host operating system exploits,” in *28th USENIX Security Symposium (USENIX Security 19)*, 2019, pp. 1357–1374.

[36] J. Lu, H. Li, C. Liu, L. Li, and K. Cheng, “Detecting missing-permission-check vulnerabilities in distributed cloud systems,” in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 2022, pp. 2145–2158.

[37] E. Pauley, R. Sheatsley, B. Hoak, Q. Burke, Y. Beugin, and P. McDaniel, “Measuring and mitigating the risk of ip reuse on public clouds,” in *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2022, pp. 558–575.

[38] P. T. J. Kon, S. Kamali, J. Pei, D. Barradas, A. Chen, M. Sherr, and M. Yung, “{SpotProxy}: Rediscovering the cloud for censorship circumvention,” in *33rd USENIX Security Symposium (USENIX Security 24)*, 2024, pp. 2653–2670.

[39] J. Niu, W. Peng, X. Zhang, and Y. Zhang, “Narrator: Secure and practical state continuity for trusted execution in the cloud,” in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 2022, pp. 2385–2399.

[40] X. Ge, H.-C. Kuo, and W. Cui, “Hecate: Lifting and shifting on-premises workloads to an untrusted cloud,” in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 2022, pp. 1231–1242.

[41] H. Xia, D. Zhang, W. Liu, I. Haller, B. Sherwin, and D. Chisnall, “A secret-free hypervisor: Rethinking isolation in the age of speculative vulnerabilities,” in *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2022, pp. 370–385.

[42] L. Cao, L. Meng, D. Stefan, and E. Fernandes, “Stateful least privilege authorization for the cloud,” in *33rd USENIX Security Symposium (USENIX Security 24)*, 2024, pp. 3477–3494.

[43] S. M. Saleh, I. M. Sayem, N. Madhavji, and J. Steinbacher, “Advancing software security and reliability in cloud platforms through ai-based anomaly detection,” in *Proceedings of the 2024 on Cloud Computing Security Workshop*, 2024, pp. 43–52.

[44] J. W. Clark, P. Snyder, D. McCoy, and C. Kanich, “” i saw images i didn’t even know i had” understanding user perceptions of cloud storage privacy,” in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 2015, pp. 1641–1644.

[45] M. Tabassum, T. Kosinski, and H. R. Lipford, “” i don’t own the data”: End user perceptions of smart home device data practices and risks,” in *Fifteenth symposium on usable privacy and security (SOUPS 2019)*, 2019, pp. 435–450.

[46] L. L. Visinescu, O. Azogu, S. D. Ryan, Y. A. Wu, and D. J. Kim, “Better safe than sorry: A study of investigating individuals’ protection of privacy in the use of storage as a cloud computing service,” *International Journal of Human–Computer Interaction*, vol. 32, no. 11, pp. 885–900, 2016.

[47] D. Wermke, N. Huaman, C. Stransky, N. Busch, Y. Acar, and S. Fahl, “Cloudy with a chance of misconceptions: exploring users’ perceptions and expectations of security and privacy in cloud office suites,” in *Sixteenth Symposium on Usable Privacy and Security (SOUPS 2020)*, 2020, pp. 359–377.

[48] K. M. Ramokapane, A. Rashid, and J. M. Such, “{“I} feel stupid i {can’t}{delete...”}: A study of {Users’} cloud deletion practices and coping strategies,” in *Thirteenth symposium on usable privacy and security (SOUPS 2017)*, 2017, pp. 241–256.

[49] E. Alhelali, K. M. Ramokapane, and J. Such, “Multiuser privacy and security conflicts in the cloud,” in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 2023, pp. 1–16.

[50] S. Kraemer and P. Carayon, “Human errors and violations in computer and information security: The viewpoint of network administrators and security specialists,” *Applied ergonomics*, vol. 38, no. 2, pp. 143–154, 2007.

[51] F. Li, L. Rogers, A. Mathur, N. Malkin, and M. Chetty, “Keepers of the machines: Examining how system administrators manage software updates for multiple machines,” in *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*, 2019, pp. 273–288.

[52] M. Kaur, H. Sri Ramulu, Y. Acar, and T. Fiebig, “” oh yes! over-preparing for meetings is my jam:)”: The gendered experiences of system administrators,” *Proceedings of the ACM on Human-Computer Interaction*, vol. 7, no. CSCW1, pp. 1–38, 2023.

[53] K. Krombholz, W. Mayer, M. Schmiedecker, and E. Weippl, “” i have no idea what i’m doing”-on the usability of deploying {HTTPS},” in *26th USENIX Security Symposium (USENIX Security 17)*, 2017, pp. 1339–1356.

[54] K. Krombholz, K. Busse, K. Pfeffer, M. Smith, and E. Von Zezschwitz, “” if https were secure, i wouldn’t need 2fa”-end user and administrator mental models of https,” in *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2019, pp. 246–263.

[55] S. Fahl, Y. Acar, H. Perl, and M. Smith, “Why eve and mallory (also) love webmasters: A study on the root causes of ssl misconfigurations,”

in *Proceedings of the 9th ACM symposium on Information, computer and communications security*, 2014, pp. 507–512.

[56] M. Bernhard, J. Sharman, C. Z. Acemyan, P. Kortum, D. S. Wallach, and J. A. Halderman, "On the usability of https deployment," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 2019, pp. 1–10.

[57] D. Nam, A. Macvean, B. A. Myers, and B. Vasilescu, "Understanding documentation use through log analysis: A case study of four cloud services," in *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 2024, pp. 1–17.

[58] M. Nassif and M. P. Robillard, "A field study of developer documentation format," in *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*, 2023, pp. 1–7.

[59] F. Fischer, Y. Stachelscheid, and J. Grossklags, "The effect of google search on software security: Unobtrusive security interventions via content re-ranking," in *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, 2021, pp. 3070–3084.

[60] M. T. Baldassarre, D. Caivano, G. Dimauro, E. Gentile, and G. Visaggio, "Cloud computing for education: a systematic mapping study," *IEEE transactions on education*, vol. 61, no. 3, pp. 234–244, 2018.

[61] D. G. Hrebec and M. Stiber, "A survey of system administrator mental models and situation awareness," in *Proceedings of the 2001 ACM SIGCPR conference on Computer personnel research*, 2001, pp. 166–172.

[62] C. Weir, I. Becker, and L. Blair, "A passion for security: Intervening to help software developers," in *2021 IEEE/ACM 43rd International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*. IEEE, 2021, pp. 21–30.

[63] I. Ryan, U. Roedig, and K.-J. Stol, "Measuring secure coding practice and culture: A finger pointing at the moon is not the moon," in *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*. IEEE, 2023, pp. 1622–1634.

[64] T. Li, E. Louie, L. Dabbish, and J. I. Hong, "How developers talk about personal data and what it means for user privacy: A case study of a developer forum on reddit," *Proceedings of the ACM on Human-Computer Interaction*, vol. 4, no. CSCW3, pp. 1–28, 2021.

[65] M. Allamanis and C. Sutton, "Why, when, and what: analyzing stack overflow questions by topic, type, and code," in *2013 10th Working conference on mining software repositories (MSR)*. IEEE, 2013, pp. 53–56.

[66] A. Barua, S. W. Thomas, and A. E. Hassan, "What are developers talking about? an analysis of topics and trends in stack overflow," *Empirical software engineering*, vol. 19, pp. 619–654, 2014.

[67] C. Parnin, C. Treude, L. Grammel, and M.-A. Storey, "Crowd documentation: Exploring the coverage and the dynamics of api discussions on stack overflow," *Georgia Institute of Technology, Tech. Rep*, vol. 11, 2012.

[68] R. Croft, Y. Xie, M. Zahedi, M. A. Babar, and C. Treude, "An empirical study of developers' discussions about security challenges of different programming languages," *Empirical Software Engineering*, vol. 27, pp. 1–52, 2022.

[69] H. Kaur, S. Amft, D. Votipka, Y. Acar, and S. Fahl, "Where to recruit for security development studies: Comparing six software developer samples," in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 4041–4058.

[70] M. Tahaei and K. Vaniea, "A survey on developer-centred security," in *2019 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, 2019, pp. 129–138.

[71] S. Kabir, D. N. Udo-Imeh, B. Kou, and T. Zhang, "Is stack overflow obsolete? an empirical study of the characteristics of chatgpt answers to stack overflow questions," in *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 2024, pp. 1–17.

[72] D. Greene and K. Shilton, "Platform privacies: Governance, collaboration, and the different meanings of "privacy" in ios and android development," *new media & society*, vol. 20, no. 4, pp. 1640–1657, 2018.

[73] K. Shilton and D. Greene, "Linking platforms, practices, and developer ethics: Levers for privacy discourse in mobile application development," *Journal of Business Ethics*, vol. 155, pp. 131–146, 2019.

[74] M. Tahaei, K. Vaniea, and N. Saphra, "Understanding privacy-related questions on stack overflow," in *Proceedings of the 2020 CHI conference on human factors in computing systems*, 2020, pp. 1–14.

[75] X.-L. Yang, D. Lo, X. Xia, Z.-Y. Wan, and J.-L. Sun, "What security questions do developers ask? a large-scale study of stack overflow posts," *Journal of Computer Science and Technology*, vol. 31, pp. 910–924, 2016.

[76] M. Tahaei, T. Li, and K. Vaniea, "Understanding privacy-related advice on stack overflow." *Proc. Priv. Enhancing Technol.*, vol. 2022, no. 2, pp. 114–131, 2022.

[77] A. Jallow, M. Schilling, M. Backes, and S. Bugiel, "Measuring the effects of stack overflow code snippet evolution on open-source software security," in *IEEE Symposium on Security and Privacy (SP)*, 2024.

[78] A. Jallow and S. Bugiel, "Stack overflow meets replication: Security research amid evolving code snippets (extended version)," *arXiv preprint arXiv:2501.16948*, 2025.

[79] S. Overflow, "Stack overflow developer survey," 2025, accessed: 2025-03-13. [Online]. Available: https://insights.stackoverflow.com/survey

[80] ——, "2023 developer survey," https://survey.stackoverflow.co/2023/, accessed: 2024-01-26.

[81] G. BigQuery, "Stack overflow dataset," https://console.cloud.google.com/marketplace/product/stack-exchange/stack-overflow, Accessed: 2023-09-01.

[82] Stack Exchange, Inc., "Stack exchange data dump," 2025, accessed: 2025-03-18. [Online]. Available: https://archive.org/details/stackexchange

[83] Internet Archive, "Internet archive: Digital library of free & borrowable books, movies, music & wayback machine," 2025, accessed: 2025-03-18. [Online]. Available: https://archive.org

[84] C. Byun, P. Vasicek, and K. Seppi, "Dispensing with humans in human-computer interaction research," in *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*, 2023, pp. 1–26.

[85] C. Shah, R. W. White, R. Andersen, G. Buscher, S. Counts, S. S. S. Das, A. Montazer, S. Manivannan, J. Neville, X. Ni *et al.*, "Using large language models to generate, validate, and apply user intent taxonomies," *arXiv preprint arXiv:2309.13063*, 2023.

[86] P. Hämäläinen, M. Tavast, and A. Kunnari, "Evaluating large language models in generating synthetic hci research data: a case study," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 2023, pp. 1–19.

[87] G. Faggioli, L. Dietz, C. L. Clarke, G. Demartini, M. Hagen, C. Hauff, N. Kando, E. Kanoulas, M. Potthast, B. Stein *et al.*, "Perspectives on large language models for relevance judgment," in *Proceedings of the 2023 ACM SIGIR International Conference on Theory of Information Retrieval*, 2023, pp. 39–50.

[88] H. Jung, W. Seo, S. Song, and S. Na, "Toward value scenario generation through large language models," in *Companion Publication of the 2023 Conference on Computer Supported Cooperative Work and Social Computing*, 2023, pp. 212–220.

[89] OpenAI, "Gpt-4," 2023, accessed: 2025-03-13. [Online]. Available: https://openai.com/index/gpt-4/

[90] M. Hogan, F. Liu, A. Sokol, and J. Tong, "Nist cloud computing standards roadmap," *NIST Special Publication*, vol. 35, pp. 6–11, 2011.

[91] F. Liu, J. Tong, J. Mao, R. Bohn, J. Messina, L. Badger, D. Leaf *et al.*, "Nist cloud computing reference architecture," *NIST special publication*, vol. 500, no. 2011, pp. 1–28, 2011.

[92] C. Rosen and E. Shihab, "What are mobile developers asking about? a large scale study using stack overflow," *Empirical Software Engineering*, vol. 21, pp. 1192–1223, 2016.

[93] G. Uddin, F. Sabir, Y.-G. Guéhéneuc, O. Alam, and F. Khomh, "An empirical study of iot topics in iot developer discussions on stack overflow," *Empirical Software Engineering*, vol. 26, pp. 1–45, 2021.

[94] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," *Journal of machine Learning research*, vol. 3, no. Jan, pp. 993–1022, 2003.

[95] S. Beyer and M. Pinzger, "A manual categorization of android app development issues on stack overflow," in *2014 IEEE International Conference on Software Maintenance and Evolution*. IEEE, 2014, pp. 531–535.

[96] C. Treude, O. Barzilay, and M.-A. Storey, "How do programmers ask and answer questions on the web?(nier track)," in *Proceedings of the 33rd international conference on software engineering*, 2011, pp. 804–807.

[97] R. Řehůřek and P. Sojka, "Software Framework for Topic Modelling with Large Corpora," in *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*. Valletta, Malta: ELRA, May 2010, pp. 45–50.

[98] B. Shao and J. Yan, “Recommending answerers for stack overflow with lda model,” in *proceedings of the 12th Chinese conference on computer supported cooperative work and social computing*, 2017, pp. 80–86.

[99] K. Charmaz, *Constructing grounded theory*. sage, 2014.

[100] J. M. Corbin and A. Strauss, “Grounded theory research: Procedures, canons, and evaluative criteria,” *Qualitative sociology*, vol. 13, no. 1, pp. 3–21, 1990.

[101] A. Strauss and J. M. Corbin, *Grounded theory in practice*. Sage, 1997.

[102] V. Braun and V. Clarke, “Using thematic analysis in psychology,” vol. 3, no. 2, pp. 77–101, tex.ids: braunUsingThematicAnalysis2006 tex.publisher: Taylor & Francis.

[103] M. M. Hennink, B. N. Kaiser, and V. C. Marconi, “Code saturation versus meaning saturation: how many interviews are enough?” *Qualitative health research*, vol. 27, no. 4, pp. 591–608, 2017.

[104] M. Hennink and B. N. Kaiser, “Sample sizes for saturation in qualitative research: A systematic review of empirical tests,” *Social science & medicine*, vol. 292, p. 114523, 2022.

[105] S. O. H. Center, “What does it mean when an answer is ”accepted”? - Help Center — stackoverflow.com,” https://stackoverflow.com/help/accepted-answer, [Accessed 08-02-2024].

[106] K. Fu, E. Sit, K. Smith, and N. Feamster, “The dos and don’ts of client authentication on the web,” in *10th USENIX Security Symposium (USENIX Security 01)*, 2001.

[107] J. Bonneau, C. Herley, P. C. Van Oorschot, and F. Stajano, “The quest to replace passwords: A framework for comparative evaluation of web authentication schemes,” in *2012 IEEE symposium on security and privacy*. IEEE, 2012, pp. 553–567.

[108] L. Geierhaas, A.-M. Ortloff, M. Smith, and A. Naiakshina, “{Let’s} hash: Helping developers with password security,” in *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*, 2022, pp. 503–522.

[109] A. Naiakshina, A. Danilova, C. Tiefenau, M. Herzog, S. Dechand, and M. Smith, “Why do developers get password storage wrong? a qualitative usability study,” in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 2017, pp. 311–328.

[110] A. Naiakshina, A. Danilova, E. Gerlitz, E. Von Zezschwitz, and M. Smith, “” if you want, i can store the encrypted password” a password-storage field study with freelance developers,” in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 2019, pp. 1–12.

[111] B. Dagenais and M. P. Robillard, “Creating and evolving developer documentation: understanding the decisions of open source contributors,” in *Proceedings of the eighteenth ACM SIGSOFT international symposium on Foundations of software engineering*, 2010, pp. 127–136.

[112] “Microsoft certified: Azure security engineer associate,” https://learn.microsoft.com/en-us/credentials/certifications/azure-security-engineer/, 2025, accessed: 2025-10-27.

[113] “Aws certified security – specialty,” https://aws.amazon.com/certification/certified-security-specialty/, 2025, accessed: 2025-10-27.

[114] C. Tiefenau, M. Häring, K. Krombholz, and E. Von Zezschwitz, “Security, availability, and multiple information sources: Exploring update behavior of system administrators,” in *Sixteenth Symposium on Usable Privacy and Security (SOUPS 2020)*, 2020, pp. 239–258.

[115] A. D. Jenkins, L. Liu, M. K. Wolters, and K. Vaniea, “Not as easy as just update: Survey of system administrators and patching behaviours,” in *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 2024, pp. 1–17.

[116] Google Cloud, “Details of the google cloud vmware engine (gcve) incident,” https://cloud.google.com/blog/products/infrastructure/details-of-google-cloud-gcve-incident, March 2024, accessed: 2025-03-27. [Online]. Available: https://cloud.google.com/blog/products/infrastructure/details-of-google-cloud-gcve-incident

[117] D. Gotterbarn, “Acm code of ethics and professional conduct,” 2018.

[118] D. Dittrich and E. Kenneally, “The menlo report: Ethical principles guiding information and communication technology research,” U.S. Department of Homeland Security, Science and Technology Directorate, Cyber Security Division, Technical Report CSD-249, Aug. 2012, accessed: 2025-08-04. [Online]. Available: https://www.dhs.gov/sites/default/files/publications/CSDMenloPrinciplesCORE20120803_1.pdf

Appendix

A. Hosting Providers

AWS	Oracle Cloud Infrastructure	Vultr
Microsoft Azure	OpenStack	Render
Google Cloud Platform	IBM Cloud or Watson	Scaleway
Firebase	Colocation	
Heroku	Cloudflare	
DigitalOcean	Vercel	
VMware	Netlify	
Managed Hosting	Hetzner	
Linode	OpenShift	
OVH	Fly.io	

TABLE III: Our list of hosting providers that we used to obtain our dataset

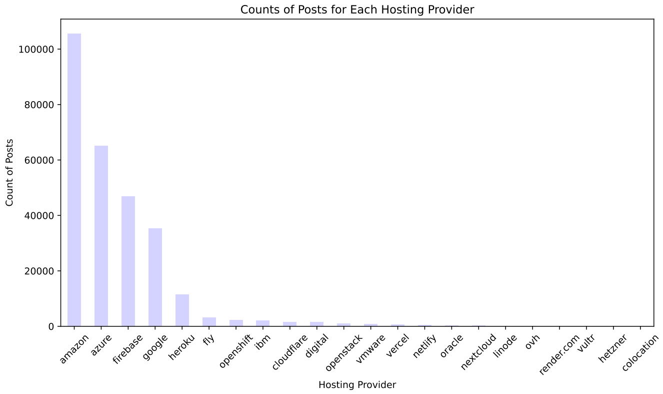

Fig. 3: Distribution of Posts across our shortlisted Cloud Providers.

B. Prompts to obtain S&P keywords

Prompts
Generate a list of security and privacy keywords for cloud resource management.
Describe the pipeline for configuring and maintaining cloud platforms like GCP or AWS, focusing on security and privacy aspects.
List security and privacy use cases, maintenance steps, defensive mechanisms, and threat models for cloud resources.
Provide keywords related to security and privacy.
Identify keywords for Stack Overflow posts concerning security or privacy.
List keywords associated with security, privacy, confidentiality, integrity, authentication, availability, authorization, accountability, and non-repudiation.

TABLE IV: Prompts used to generate security/privacy keywords by querying the OpenAI ChatGPT web app with the GPT-4 model during October 27 to November 1, 2023. We did not include ”access control” in our list of prompts as our initial rounds of pre-testing led us to conclude that the terms ”authentication” and ”authorization” already obtain relevant responses.

C. Metrics to evaluate LDA

See Figures 4 and 5.

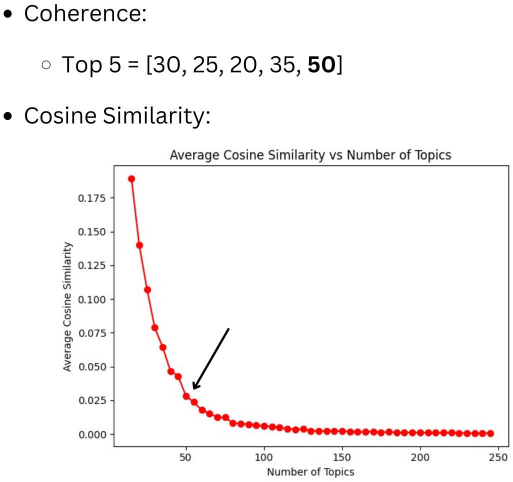

- Coherence:
  - Top 5 = [30, 25, 20, 35, **50**]
- Cosine Similarity:

Fig. 4: We used cosine similarity and coherence to determine the best number of topics. The results of these metrics for the 2008 - 2022 dataset are displayed in this figure.

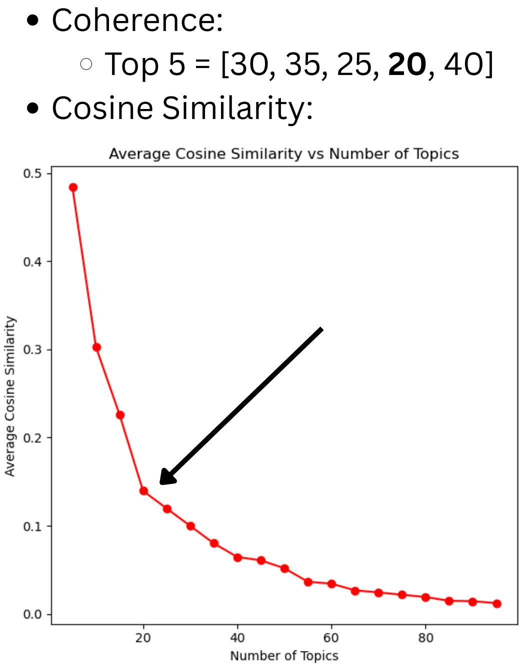

- Coherence:
  - Top 5 = [30, 35, 25, **20**, 40]
- Cosine Similarity:

Fig. 5: Cosine similarity and coherence results of these metrics for the 2022 - 2024 dataset are displayed in this figure.

D. Distribution of Posts Across Time

See Figure 6.

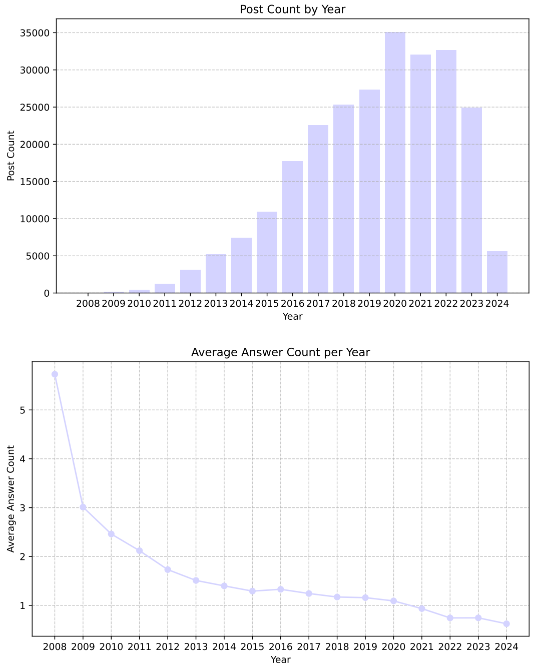

Fig. 6: Temporal plots of our Stack Overflow dataset: (a) number of questions posted each year, and (b) average answer count over the years.

E. Documentation Analysis

See Figure 7.

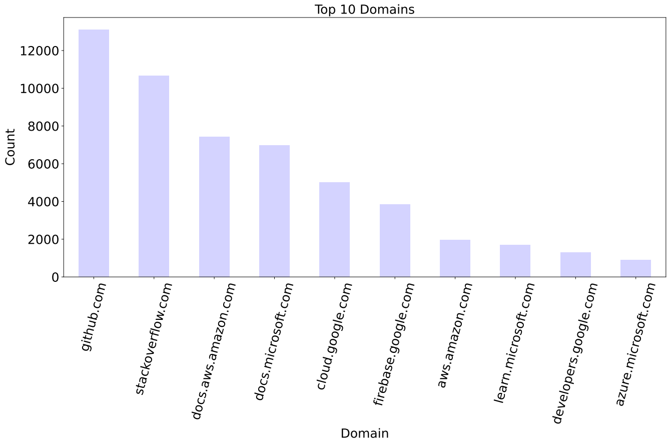

Fig. 7: Top 10 Most Frequently Cited Websites