

***“It’s almost like Frankenstein”*: Investigating the Complexities of Scientific Collaboration and Privilege Management within Research Computing Infrastructures**

Souradip Nath^{*†}, Ananta Soneji^{*†}, Jaeyong Baek^{*}, Tiffany Bao^{*}, Adam Doupé^{*},
Carlos Rubio-Medrano[‡], and Gail-Joon Ahn^{*}

^{*}Arizona State University, [‡]Texas A&M University – Corpus Christi

^{*}{snath8, asoneji, jbaek7, tbao, doupe, gahn}@asu.edu, [‡]carlos.rubiomedrano@tamucc.edu

Abstract—Research Computing Infrastructures (RCIs) integrate high-performance computing, advanced data storage solutions, and sophisticated network protocols, connecting people, data, and computing resources to facilitate scientific collaboration in today’s data-driven world. Access control is essential in such highly collaborative environments to prevent resource misutilization, safeguard data integrity, and allocate resources effectively, thereby enabling secure and trusted interactions among different users. However, unlocking the full potential of RCIs for collaborative research through effective access control requires more than technological exploration—it demands a deep, human-centered understanding of the stakeholders who operate and utilize these systems.

In this paper, we present the first qualitative study that explores the human dimensions of RCI interactions, drawing insights from 24 key stakeholders, including researchers and system administrators, across 12 research institutions to examine the collaborative practices, challenges, and needs with a focus on access control. Our findings reveal operational complexities and project-specific, trust-based resource-sharing dynamics, highlighting tensions between security and usability. Based on these insights, we provide stakeholder-driven recommendations and requirements for adaptive, user-centered access control for RCIs, laying the groundwork for advancing human-centered security practices in RCIs.

1. Introduction

In the 21st century, scientific research has entered a new era, where advancements in computational and data resources drive discovery and innovation across disciplines. Recognized by the National Science Foundation (NSF) as the “fourth paradigm” of science, data-driven research—alongside experimental, theoretical, and computational approaches—is fundamental to advancing scientific knowledge [1]. This shift has driven a critical demand for high-performance computing (HPC), enabling researchers to derive insights, make predictions, and support complex decisions in science and engineering [2], [3].

HPC has long been essential for computational research, helping researchers to tackle complex, data-intensive problems across fields ranging from training large language models [4] to genetic data sequencing [5], [6]. Recent advancements in cloud computing and scalable cluster technologies have allowed HPC resources to integrate seamlessly into scientific cyberinfrastructures, commonly referred to as Research Computing Infrastructures (RCIs). These environments leverage the computational power of HPC, along with vast data storage (ranging from terabytes to petabytes) and high-speed networking (up to 200 Gbps), creating an ecosystem that brings people, information, and computational tools together to support all aspects of research computing [7].

Consequently, academic research institutions are increasingly adopting RCIs, consolidating resources that were previously managed in siloed research environments by different research groups. Recent data from the TOP500 list, which ranks the world’s most powerful supercomputers based on performance, reveal that many of the top HPC centers are part of RCIs managed by academic research institutions, demonstrating a growing role of RCIs in advancing scientific research within academic disciplines [8].

As scientific research grows more collaborative, RCIs facilitate security-sensitive resource sharing across complex collaboration workflows. Unauthorized access to RCI resources poses serious risks, including data breaches and resource misutilization, such as the cyberattack on the Framingham Heart Study (FHS) at Boston University [9], where unauthorized individuals accessed sensitive data from more than 15,000 participants. This underscores the critical necessity for robust access control measures to maintain the overall security posture of RCIs [10], [11].

Although advancements in access control offer essential security measures for these infrastructures (§ 2), a solely technical approach may miss important nuances in stakeholder interactions that affect secure and efficient resource sharing. Additionally, as RCIs expand, a holistic approach is needed—one that considers the dynamics of those who operate and rely on these systems—to effectively safeguard these infrastructures.

In light of these challenges, we conduct the first quali-

[†]. These authors contributed equally to this work

tative investigation focusing on the human aspects of secure collaboration, resource sharing, and privilege management within RCIs. Through in-depth, semi-structured interviews with 24 key RCI stakeholders, including researchers and system administrators, across RCIs in 12 distinct institutions, we aim to uncover their practices, challenges, and needs around secure scientific collaboration and access management. This dual focus on access control and collaborative research practices positions our work at the intersection of cybersecurity and computer-supported cooperative work (CSCW), addressing the following research questions:

RQ₁: How is collaborative research enabled in RCI ecosystems at research institutions, including architectural and procedural aspects?

RQ₂: How are collaborative research activities accomplished within RCI ecosystems?

RQ₃: What technical and non-technical challenges do current RCI stakeholders experience, and what improvements do they envision for the future?

Based on our analysis, we present several novel observations on the operational and collaboration dynamic within RCIs highlighting key findings:

Procedural Dynamics of RCI Management: Collaborative research within RCIs is shaped by a combination of user roles, resource ownership dynamics, and delegated privileges that create a nuanced, layered structure for access control. This interplay of shared responsibilities challenges traditional access control, demanding adaptive approaches to support secure, collaborative workflows.

Collaboration Dynamics and RCI Utility: Researchers frequently form teams driven by both resource demands and interdisciplinary goals, positioning RCIs as enablers of collaborative scientific work. Our study also reveals that informal, trust-based resource-sharing practices often coexist with formal access control protocols, though this raises security compliance concerns.

Critical Challenges of RCI Stakeholders: RCI stakeholders face diverse technical and organizational challenges, including limited automation, fragmented infrastructure, coarse-grained access control, and inadequate support systems. While these issues complicate secure and efficient access control, they also impact usability and collaboration for researchers. Despite these obstacles, sentiment analysis reveals that stakeholders value RCIs but express a need for improved training, support, and shared autonomy to enhance system usability with robust security.

Contributions. Our study makes the following contributions:

- We present the first human-centered qualitative study of RCIs, offering a nuanced understanding of the practices, challenges, and collaborative dynamics within these infrastructures. By prioritizing the human perspective, our study provides foundational insights often overlooked in purely technical assessments of RCIs.
- By capturing experiences from both researchers and administrators, we provide a unique multi-perspective analysis of RCIs, revealing how they navigate, manage, and

perceive these infrastructures. This approach sheds light on critical themes around ownership dynamics in privilege management, informal trust-based practices, and security-usability tensions.

- Based on these human-centered insights, we offer actionable, stakeholder-informed recommendations to enhance security, usability, and system interaction in RCIs. Our work also proposes essential requirements for designing secure, adaptive access control frameworks that support evolving research needs within RCIs.

2. Background and Related Work

Understanding RCIs requires both technical knowledge and an appreciation for the interactions among the human stakeholders who design, maintain, and rely on these infrastructures. This section outlines the technical overview and key stakeholders of RCIs, followed by a review of relevant literature on security and human factors in RCI.

Technical Overview. RCI resources can broadly be classified into two categories, namely, *computing resources* and *data directories*. At the core of RCIs are *supercomputing clusters*, built from thousands of high-performance CPU cores and hundreds of GPU accelerators, organized into partitions that collectively form a cluster [12], [13], [14], [15], [16], [17]. Such architecture enables massive parallel processing, allowing researchers to scale computationally intensive tasks such as data analytics and scientific modeling.

In addition to computing power, RCIs offer data storage, from high-speed temporary storage for jobs to project-specific long-term repositories and archival storage to manage and share large-scale datasets. High-speed networks, typically powered by InfiniBand (IB), ensure connectivity across computing, storage, and research facilities [18].

Resource allocation and access control around computing resources are typically managed by job schedulers, such as Slurm [19], in conjunction with access control systems such as Pluggable Authentication Module (PAM) [20]. Slurm-like job schedulers are designed to allocate resources dynamically based on job priorities, user quotas, and system load, while PAM in Unix-based systems, provides a unified approach to authentication and session management, ensuring that only authorized users can access specific resources. Together they enforce flexible access policies while dynamically scheduling and regulating resource access. On the other hand, access privileges around data directories are managed using POSIX permissions [21].

Human Stakeholders. Academic and research institutions are the primary adopters and users of RCIs, driving the demand for these high-performance resources [22]. System administrators, including HPC architects, network engineers, and cybersecurity experts, orchestrate, maintain, and manage the entire infrastructure, along with designing access control mechanisms within RCIs, meeting research needs and supporting secure collaboration. Researchers, including faculty members (often termed Principal Investigators (PIs)), graduate students, and doctoral and postdoctoral researchers,

are the primary users of RCIs, relying on these resources for complex, and highly collaborative research [23].

Security in RCI. As RCIs have emerged, security and access control challenges within RCIs have consequently scaled with their expansion [10], [11]. Recent research on RCI security has emphasized the critical need for secure, scalable collaboration and dynamic resource sharing mechanisms. Studies have introduced granular, context-aware data-sharing frameworks to meet the demands of multi-institutional research initiatives, implementing novel models for secure authentication and data access management [24], [25], [26], [27], [28]. Further, federated identity solutions have become central to supporting secure, cross-domain research collaborations, allowing institutions to streamline user access without compromising security [29]. A notable focus has also been on enhancing the interoperability and adaptability of access management approaches, addressing the need for flexible, scalable control over resource access [27], [30], [31]. Recent work has leveraged AI to proactively identify and mitigate risks within HPC systems [32].

Human Factors in Security and RCI. Human factors research in security demonstrates how social behaviors [33] and influence [34], different organizational roles [35], expertise-levels [36], and decision-making processes [37] shape security practices. Studies using qualitative methods, e.g. interviews, have proven essential in understanding human behaviors, emotions, practices, and experiences within cybersecurity contexts [38]. For instance, research has examined the complexities of deploying intrusion detection systems [39], reviewing access control policies [40], evaluating security research [41], and addressing real-world practices in vulnerability management [42], highlighting the critical role of human dynamics in security.

While studies in RCI have begun to touch on the human dimensions, the role of human factors remains an emerging area. RCIs rely not only on advanced technology but also on the effective collaboration and social practices among stakeholders. Earlier studies introduced the concept of “human infrastructure” [43] within large cyberinfrastructure projects, showing that successful collaboration depends on synergy among team members as much as on technical systems [44]. Although more recent studies [45], [46] have engaged RCI stakeholders in user-centered research, there is still a notable gap in integrating human perceptions, practices, and needs into security solutions for RCIs.

Building on this body of work, our research investigates the interactions of two primary RCI stakeholders—researchers and admins—for resource sharing and privilege management. By exploring their practices and challenges through interviews, we seek to provide a foundational understanding of how human factors influence access control and user experience within RCIs. Through our findings, we also attempt to offer actionable recommendations and inspire future research that fosters a more secure and user-centric environment for all RCI stakeholders.

3. Methodology

To gain a comprehensive understanding of the scientific collaboration and access control practices, challenges, and needs within RCIs, we conducted semi-structured interviews with 24 RCI stakeholders from 11 research universities and 1 national laboratory. Our participants represented two primary groups: *Researchers*—comprising Principal Investigators (PIs) and non-PI (postdoctoral researchers and graduate students)—and *System Administrators*. Researchers provided insights into computational needs and collaboration dynamics, while administrators offered perspectives on operational complexities. This selection allowed us to capture perspectives from key stakeholders who interact with RCI in complementary but distinct ways. Overall, we interviewed 14 RCI researchers (denoted as R01 – R14), i.e., six PI and eight non-PI researchers, and 10 administrators (denoted as A01 – A10).

3.1. Participant Recruitment

To recruit a diverse sample of RCI admins and researchers, specifically targeting hard-to-reach groups—PIs and admins [47], we employed a multi-faceted approach:

Initial Participant Identification: We compiled a list of universities with large-scale RCI by reviewing university websites and academic databases. From each RCI site, we gathered contact information for system administrators and identified PIs involved in research using RCI resources. We then expanded our approach by identifying non-PI collaborators listed on the PIs’ professional web pages, establishing an initial contact list of potential participants.

Social and Professional Networks: To enhance our outreach and quickly disseminate study information, including objectives and participant eligibility criteria, we posted recruitment announcements on X (formerly known as Twitter) and LinkedIn. Additionally, we leveraged our professional networks within the research community to directly invite eligible researcher participants.

Snowball Sampling Technique: We employed snowball sampling technique [48], effective for accessing hard-to-reach populations [47], to recruit admins. Given their close-knit professional communities, referrals from admin participants facilitated connections with additional RCI admins.

Of our 24 study participants, 12 were recruited through direct email outreach, five via snowball sampling, four through social media, and three via professional network.

Screening Survey and Eligibility Criteria: Our recruitment approach, including email invitations and social media outreach, incorporated a screening survey (see Replication Package [49]) that interested participants completed. Participants were then selected based on defined eligibility criteria. Based on their survey responses, we invited *Researcher* participants who had (a) at least 1-2 years of cumulative experience actively engaging with a large-scale RCI for research purposes, and (b) demonstrate active collaboration with a minimum of 4–5 peers sharing these resources. For

System Administrator participants, we invited those with a minimum 1 year of experience actively managing a large-scale RCI.

We defined ‘large-scale’ RCIs based on metrics such as the number of clusters (2 or more), computing nodes (hundreds or more), CPU cores (in thousands), GPUs (in hundreds), storage capacity (at least one petabyte), and institution size (R1 and R2 institutions, as per the Carnegie Classification¹).

Our social media campaign received 16 responses, but only 4 were selected after a rigorous vetting process. This included strict adherence of eligibility criteria, thorough background checks through public profiles (e.g., professional websites, publications, and work history), and follow-up emails to clarify information regarding participants’ current affiliations and their RCI experience. During this process, we recruited an additional researcher affiliated with a national laboratory, after confirming their substantial prior RC experience within a university setting and current professional role. In general, any identified discrepancies or concerns regarding participant quality resulted in exclusion.

Overall, we conducted interviews with 24 eligible participants out of 47 interested participants. Among the remaining 23 individuals, 14 were excluded as a result of the vetting process, 6 due to profile similarity with existing participants to balance diversity, and 3 withdrew after filling out the screening form. Lastly, our multi-faceted recruitment approach and rigorous vetting process resulted in multiple participants from 4 out of 12 institutions, with three institutions represented by both researcher and admin participants.

3.2. Interview Design

We developed two distinct sets of interview questionnaires tailored to the roles, and experiences of two key stakeholder groups: *Researchers* and *System Administrators* (See Appendices A and B). Both questionnaires included overlapping questions to capture commonalities and differences in perspectives, enhancing our understanding of collaboration and access control dynamics. Overall, our interview design, illustrated in Figure 1, included distinct sections, each introduced with a brief overview of the topics covered. Researcher questionnaire aimed to explore their motivations for collaboration (1a, Fig: 1), practices related to resource-sharing (2a, Fig: 1), challenges encountered in collaborative research environments, and potential needs regarding access management within their teams (3, Fig: 1).

Conversely, admin questionnaire focused on gaining insights into the current RCI landscape (1b, Fig: 1), the access control mechanisms, and the user and resource management practices (2b, Fig: 1). Additionally, we inquired about the desired features for access control from the administrator’s perspective (3, Fig: 1), highlighting their role in shaping access management practices.

Piloting: Before finalizing the questionnaires, we pre-tested them with three pilot interviews—one admin, one PI,

1. Carnegie Classification of Higher Education Institutions: <https://carnegieclassifications.acenet.edu/>




Figure 1: **Interview Protocol.** Each interview section began with broad, open-ended questions designed to explore participants’ general experiences with RCI. Certain sections were uniquely tailored to reflect the specific roles and backgrounds of each participant group within RCI, with Sections 1a and 2a focusing on Researcher practices, and Sections 1b and 2b addressing Administrator practices. Other sections, such as Sections 3 and 4, covered themes common to both groups, including shared challenges, needs, and overall satisfaction with their RCI experiences.

and one student researcher. Data from these pilot interviews is not included in our final results.

Pilot interviews were instrumental in enhancing the clarity, neutrality, and overall effectiveness of our questionnaire. They enabled us to identify and revise questions that could introduce bias or limit open discussion. For instance, Q7 (Appendix A) was changed from “Do you or your team follow specific rules on access and resource sharing?”—which was closed-ended—to “In your collaborative teams, how is it decided who is going to access what resources? Are there any specific rules or policies in place?”—for richer and more candid participant insights.

3.3. Data Collection

To manage data collection efficiently, the primary researcher sent interview invitations—detailing the study’s motivation, procedures, ethical considerations, compensation details, and the screening survey—in small batches and conducted all interviews remotely via Zoom from August

Academic Role	#
Assistant Professor	5
Professor	1
Postdoctoral Researcher	2
Graduate Student	6
Research Discipline	#
Computer Science and Other Engineering	5
Biological and Biomedical Sciences	5
Earth and Geographical Sciences	3
Other (e.g., Cognitive Science)	1
Experience with Research Computing	#
less than 5 years	7
6 to 10 years	4
more than 10 years	3
Approximate Number of Active Collaborators	#
less than 10	5
11 to 20	4
more than 20	5
Familiarity with Access Control	#
Not at all familiar	4
Slightly familiar (fundamental awareness and basic knowledge)	4
Somewhat familiar (limited experience)	6
Total	14

TABLE 1: **Researcher Participant Demographics.** An overview of the participants, detailing their academic roles, research domains, experience with collaborative Research Computing, and familiarity with Access Control practices.

2023 to June 2024. Interview durations varied, ranging from 27 to 80 minutes, with an average length of 52 minutes.

At the beginning of the interview, we confirmed participants’ consent to be recorded. We used a semi-structured interviewing format with open-ended, non-leading questions which allowed for free and flexible conversation with our study participants. With this format, participants could skip questions or request clarification, helping create a comfortable environment to share insights. When necessary, the interviewer could ask follow-up questions to explore specific areas further or address points not fully covered in participants’ responses. Additionally, prior to the interviews, participants were asked to complete a demographic survey [49], which helped contextualize their responses during the interviews.

Several key demographic details of our participants are presented in Tables 1 and 2. The *researcher* participants included both PI and non-PI researchers from diverse fields such as Computer Science, Biomedical Sciences, and Earth Sciences. While all researchers had substantial experience in research computing, their familiarity with access control was generally limited, spanning minimal to basic levels of knowledge. In contrast, *administrator* participants primarily came from Information Technology (IT) and Cybersecurity backgrounds, with considerable experience in RCI administration and a more advanced, applied understanding of access control.

Educational Background	#
Computer Science	3
IT/Cybersecurity	6
Other (e.g., Physics, Mathematics)	1
Experience with RC Administration	#
less than 5 years	6
more than 5 years	4
Familiarity with Access Control	#
Somewhat familiar (limited experience)	1
Moderately familiar (applied theory and practical knowledge)	8
Extremely familiar (recognized authority)	1
Total	10

TABLE 2: **Administrator Participant Demographics.** An overview of the participants, detailing their educational background, experience with Research Computing administration, and familiarity with Access Control.

3.4. Data Analysis

We used MAXQDA software [50] to perform thematic analysis [51] on anonymized interview transcripts.

Two primary authors independently open-coded (i.e. double-coded [52]) six interviews to create an initial codebook with broader categories. The preliminary codebook was then independently applied to code eight additional transcripts, resulting in new open codes. After collaborative discussions and further categorization, a refined codebook emerged, comprising 14 overarching themes [49] based on 14 interview transcripts.

The refined codebook’s applicability was tested by independently double-coding four new transcripts, resulting in an inter-coder reliability score of 0.86 for Cohen’s Kappa ($\kappa > 0.8$: almost perfect agreement) [53]. Given this high inter-coder reliability, the codebook was finalized without requiring major thematic revisions. Consequently, the final six transcripts were analyzed individually by dividing them between the two coders, using the finalized codebook.

We determined our sample size ($n = 24$) based on theme saturation and diversity considerations. By the 18th interview, we achieved high inter-coder reliability with no significant changes to the codebook at the theme level, indicating thematic saturation. To account for potential diversity-related variations among participants (e.g., institutional affiliation, experience levels), we continued interviewing until confidently reaching saturation at 24 participants.

Our analysis identified key themes capturing challenges and practices within RCI contexts. A complete overview of these themes, along with detailed codes and definitions, is provided in our replication package [49]. Additionally, Figure 2 visually illustrates our study design, presenting the research questions (RQs), their corresponding objectives, and the RCI stakeholders consulted to address each RQ. This visual representation is aimed to support a clearer understanding of how the identified themes from our analysis directly relate to each RQ, guiding the organization of results presented throughout the paper.




Figure 2: **Overview of Results.** This figure shows the study’s research questions, consulted RCI stakeholders, and key themes from the interview analysis, providing a snapshot of findings aligned with study goals.

3.5. Ethics

This study was reviewed and approved by the primary author’s Institutional Review Board (IRB), which granted an ‘exempt’ determination [54]. Nevertheless, we adhered to institutional ethics policies for human subjects research. All participant data was anonymized and securely stored, and participants were fully informed of their rights, including the option to withdraw at any time without consequence.

In line with ethical research practices, we compensated our study participants for their time and contributions [55]. The compensation amount of \$100 was determined based on the ‘expertise’ and ‘scarcity’ of our participants, who possess specialized technical knowledge and are less likely to engage in general interview studies—particularly hard-to-reach groups (i.e., system administrators and PIs). While this amount may exceed typical compensation for non-PI participants, e.g., graduate students, our intent was not to set a precedent for their future compensation in similar studies, but rather to ensure *equitable* recognition of *all* our participants’ time and effort.

3.6. Study Limitations and Future Directions

As with qualitative research, our study is susceptible to social desirability, confirmation, self-report, and recall biases [56]. Although dedicated efforts were made to minimize bias, such as piloting and the intention of framing the questions neutrally, certain question formulations (e.g., Q11, Q12, Appendix A) may have had unintentional limitations. However, our participant responses reflected balanced perspectives, including both positive and negative experiences (§ 6.2.4), indicating that the question phrasing did not unduly influence their answers.

Our compensation structure may have led to self-selection of participants, potentially attracting individuals with stronger opinions or greater financial incentives. However, comparable study invitation acceptance rates across participant groups—non-PI researchers (17%), admins (13%), and PIs (7%)—suggest that compensation did not disproportionately influence participation. Nonetheless, variations in participant engagement during interviews may have led to an imbalance in perspectives, particularly between more and less vocal respondents. Future studies might explore alternative recruitment strategies, such as tiered compensation models and anonymous participation, to balance engagement levels and capture a broader range of viewpoints.

While our participants (Tables 1, 2) represented varied expertise, research backgrounds, collaborations, and RCI familiarity, findings are most relevant within the studied institutions. Broader applicability would benefit from additional sampling across diverse institutional locations and regulatory contexts. Although our recruitment method led to the inclusion of one non-US system administrator, our study results do not include any location-based insights.

To capture role-specific insights, we used distinct question sets for researchers and administrators, maintaining some overlap to identify shared themes; however, future studies might benefit from focus groups to explore cross-role interactions. Although three out of 12 institutions represented by both researcher and admin participants (§ 3.1), this number was insufficient for within-institution comparisons of access control and privilege management practices. Future work could systematically examine institutional-level variations, including how policy enforcement, IT governance structures, and security cultures differ across RCIs.

Finally, this study captures a snapshot in time; as RCIs technologies and institutional policies evolve, future

studies could track shifts in stakeholder needs and challenges longitudinally.

Based on our analysis of 24 interviews with RCI researchers and system administrators, we next present the results that highlight their RCI practices and challenges. Our findings are organized according to our study RQs and key insights from the semi-structured interviews with RCI researchers and system administrators, as illustrated in Figure 2.

4. (RQ1) Procedural Dynamics of RCI Management: Users, Resources, and Privileges

Understanding how collaborative research is enabled in RCI ecosystems (RQ_1) begins with examining its architectural and procedural foundations. In our study, we asked our admin participants how they develop policies, processes, and procedures around user, resource, and access management.

4.1. User Management: Roles and Sponsorship

User management in RCIs is primarily structured around a *sponsorship* model, with little to no *technical* distinction between different types of users. Central to this structure is the requirement for students to obtain sponsorship from a PI, who serves as a “*root of trust*” within this ecosystem. As A03 mentioned, “*students need to be working with a PI in order to have access... they have to be sponsored by someone*”. This sponsorship model applies broadly across user types, encompassing internal students, external collaborators, research assistants, and others.

Along with the differences in academic status—such as undergraduate, graduate, or postdoctoral roles—and institutional affiliation, temporal factors can also potentially play a role in distinguishing users. For instance, undergraduate students are often transient, “*they join for a semester or a year*” (R11) while PIs tend to have a more permanent presence within the system.

Despite these variations, A06 noted that, on a system level, all user accounts are treated uniformly:

Within the system, there's no difference, they're all users. They're users or admins, right? The users are users, there's no hierarchy within them.

Takeaway: Within a technically equal and simplified user management in RCIs, distinctive factors such as the academic status, institutional affiliation, and temporal dynamics are not considered.

4.2. Resource Management: Ownership Dynamics

Within RCIs, managing computational and data resources follows established processes. However, while data ownership is more straightforward, ownership of computational resources is more complex and nuanced.

The Condo Model for Computational Resources. The condo model of resource management functions as a collaborative framework where researchers contribute their own physical hardware to RCI while benefiting from access to a shared pool of resources and administrative support.

This arrangement creates a mutually beneficial, win-win scenario for both administrators and researchers. As A10 highlighted, “*Great for [PIs], they don't have to administer those machines. Great for me, I have more backfill resources.*” In this flexible model, researchers can directly support the infrastructure through financial contributions, which grants them prioritized access to the resources they fund. This setup creates a direct association between resource ownership and access privileges. As A04 pointed out:

When a group or PI purchases dedicated hardware, they get a boost automatically for using their own hardware, always being at the front of the line for their resources.

Public vs. Private Resource Management. While the condo model offers clear benefits to both researchers and administrators, it also creates a nuanced and complex ownership dynamic for computational resources. Although researchers retain priority over the resources they purchase, these resources become part of a shared, public pool when not in use, as A10 explained:

We pretty much enforce that if you're not using the resources, we may, at our discretion, borrow them for interruptible jobs... We don't do anything that requires a machine to be cordoned off and can only be accessed by these people who bought it.

This approach maximizes resource availability but blurs the lines of ownership, creating a collaborative environment where personal investments benefit the broader community.

User vs. Group in Data Ownership. Ownership of data directories within RCIs is generally straightforward, as researchers are assigned personal directories for storing their datasets. However, when these directories are used in collaborative research, the need for shared access among multiple users introduces additional coordination and management. To that end, personal directories need to be configured for shared access, allowing multiple users to access data.

In addition to individually owned directories, PIs have the option to purchase customized storage solutions for their entire research team, providing shared access and enhancing flexibility. As A10 aptly explained with an example:

Users, by default, get a certain amount of the shared file systems carved out for them. If they're part of research groups, we actually have separate group directories that are owned and shared by the group. So, if user A is part of Group B, then they would get access to '/home/A' and '/groups/B'.

This arrangement allows a balance between individual ownership and group collaboration, creating a multi-layered system of ownership that transitions from user-specific control to broader group-level access as needed. However, this introduces more complex challenges such

as managing access permissions, potential data confusion, resource contention, and increased security risks (§ 6.2.1).

Takeaway: *The diverse landscape of resources in RCIs, combined with a layered and nuanced notion of ownership for both computational and data resources, blurs the boundaries of authority for access-related decision making.*

4.3. Privilege Delegation and Dependencies

Effective privilege management is essential for facilitating secure scientific research and collaboration, ensuring that access to resources is carefully regulated. Building on our insights on the complexities of ownership in resource management, we sought to clarify whether all users have equal access to resources or if certain permissions are tied to factors such as ownership, user roles, or group memberships.

Group-level Policies and Control. All admin participants emphasized the role of group-based access control in efficient resource management and collaborative data sharing. Based on our interviews, we identified several methods that admins use to group users within RCIs to manage privileges and facilitate resource sharing, including grouping by *PI*, *project*, and the *condo model*.

PI-level Grouping: The most common method (A03–A08, A10) for grouping users is at the ‘PI level.’ Researchers affiliated with the same PI are grouped together and granted access to shared resources, reflecting the assumption that they are *natural collaborators*. This approach is useful for researchers working under a single PI who need to access resources for their projects. A03 explained this process:

So, we have it all managed at the PI level or some faculty member who has people working under them... we give them their own Linux group. And from that, we create all other Slurm accounts under that PI's group.

Project-level Grouping: Another method of grouping users is at the ‘project level.’ This is used when researchers, often from different PIs, collaborate on the same project and need shared access to project-specific resources. While project-level groups can overlap with PI-based groups, they offer greater flexibility for cross-PI collaboration. A04 described this structure:

Each PI will have their own group. And then we'll do like a separate third group or supplemental group. This will allow you to have access to maybe data A but not data B because of how we applied the supplemental group permissions.

Condo Model Grouping: In the condo model, group-level access is key to managing PI-contributed resources. Here, hardware is purchased and prioritized for use by specific groups, typically defined by the PI or the research team. As A10 explained:

[Access management within condo model] is performed on a group level essentially; if a research group or their PI initiates a purchase of hardware, they say I want it to be prioritized for this group of individuals.

Delegated Privilege Management to PIs and beyond. A common theme in privilege management is the delegation of authority over resource access to the owners and users of those resources (§ 4.2). This typically, as noted by all our participant admins, starts with PIs, who serve as the “*first level controllers*” (A07). PI ownership manifests in two key ways: first, through the resources they have financially contributed to the infrastructure, and second, through the resources associated with users they sponsor under their accounts. As A05 explained, decisions about access are often pushed to PIs, as they have the best understanding of who should have access:

We're trying to push down [the access decisions] to the faculty because they obviously have the most information about what students or personnel are appropriate to have access to this.

Beyond PI-owned resources, users maintain control over their personal data directories through Discretionary Access Control (DAC) [57], allowing them to manage permissions and “*share them as they wish*” (A10). A05 shared that PIs may delegate authority over resources to trusted team members, such as postdocs or graduate students, and in some cases, privilege management becomes a collaborative effort, with project members jointly deciding on permissions. This approach reflects a natural delegation of authority, with decisions made based on the academic role hierarchy, even if not formally embedded in the user management (§ 4.1).

Our interviews revealed several reasons for this delegation. First, system administrators cannot manage fine-grained user-level privileges due to system scale, as A07 noted: “*We cannot control exactly who is doing what within their group.*” Delegating this responsibility reduces administrative burden. Second, the academic role hierarchy naturally positions PIs as sponsors, fostering trust and responsibility. As A05 explained: “*All accounts are sponsored through individual faculty members, so there's sort of a root of trust that we maintain.*” Finally, administrators are not domain experts in scientific fields, making it more practical for researchers with specific expertise to manage resource access (A07: “*We're not the domain science experts*”).

Dependence on Admins on Shared Decision Making. As discussed previously, while a shared responsibility model for access decision-making is informally followed, the implementation reveals a tension between delegation and control. In theory, PIs should manage access to the resources they contribute, but in practice, they depend on admins to enforce these decisions. This paradox arises from technical limitations and the need to balance security with administrative burden. As A03 explained:

So [the PIs] get full control of which accounts are allowed to run on a node. They don't have

any special system permissions like they're not the ones that go in and add that; it's a ticket to us. And then we review it and apply those changes.

While delegating privilege management offers benefits (as discussed earlier), admin participants highlighted concerns about misconfigurations, lack of technical expertise, and security risks necessitating centralized management:

To be honest, a lot of the faculty members don't touch these systems very much... So, if they are the ones that are configuring the access control, there's a possibility an error goes off. (A06)

These security concerns restrict the extent to which non-admin users can share the responsibility of privilege management. As A04 summed up, “it’s a gray area about how much access you give them versus how much you don’t give them.” Although some institutions have developed bespoke local solutions (A04, A05 and A09) that allow PIs and others to manage their assets through self-service portals, these solutions remain neither widely adopted nor standardized.

Ownership-agnostic Privilege Management. Beyond PI-owned and project-specific resources, general access policies in RCIs are governed by a uniform set of privileges for all users lacking further granularity or tailored considerations. As stated by A03:

The point of the cluster is that it's a free publicly available cluster for any researcher to use. So the main thing we have is that you have to be a researcher at [the institution] doing something for the benefit of [the institution].

While this coarse-grained approach fosters an inclusive environment for researchers, it may inadvertently overlook the varying needs and responsibilities of different users, further elaborated in the following sections.

Takeaway: While each stakeholder within RCIs contributes uniquely to aspects of access control, the nuanced ownership dynamics (§ 4.2) can complicate defining and formalizing the scope of this shared responsibility.

5. (RQ2) Collaboration Dynamics and RCI Utility

In this section, we examine how collaborative research activities are accomplished within RCI ecosystems (RQ₂), with a focus on collaborative research team formation, collaboration motivations, and resource-sharing practices.

5.1. Motivations and Team Structures

RCI enables research teams to collaborate effectively by facilitating complex computations and managing data. However, to fully understand the use cases of RCI for collaborative purposes, as well as the practices and challenges involved, it is essential to explore the reasons why researchers form collaborative teams.

Motivations behind Collaborative Research Teams. Participants in our study formed research teams to leverage expertise, overcome resource limitations, and pursue interdisciplinary projects. R10 emphasized the importance of collaboration by stating, “If you want to really address any significant problem, you usually need expertise from multiple specialties.” Similarly, R13 underscored the value of combining team expertise with external collaborations, often with specialized backgrounds. In addition to expertise, researchers such as R14 often collaborate due to “lack of availability of resources in the current lab,” particularly in cross-domain projects. They explained:

If I'm working on security and I want to build a machine learning project, then I contact the lab that specializes in machine learning, not security, so that we can collaborate and get the work done.

R04 acknowledged that research fields are “massive,” making collaborations essential for accessing “more concentrated knowledge of a particular domain,” particularly in interdisciplinary research areas of common interest.

Roles and Team Structures. Participants provided insights into research teams composition and structures. R02 described the primary author and the PI as “primary stakeholders in the matter.” R04 described different team roles:

- *Project leaders* originate and drive the project, often refining ideas and leading experiments.
- *Supporting authors* assist with scaling experiments or offer advice, typically without technical involvement, including senior Ph.D. students or mentors.
- *PIs* provide high-level guidance, often without direct involvement in technical tasks.

While hierarchical structures streamline responsibilities (R05) and mentorship (R07), some participants prefer ad-hoc arrangements for flexibility, particularly among well-trained members. R06 shared: “Whenever I work with other people, it’s more loose... we all have a similar goal... and we kind of just talk about what we’re doing and share code.”

R05 and R07 pointed out that team composition and role assignments are influenced by academic positions, funding, and project-specific needs, underscoring the adaptability of research teams within RCIs.

Takeaway: Researchers form collaborative teams to bridge expertise gaps and tackle interdisciplinary challenges, with clear roles and adaptive structures.

5.2. RCI Utility for Collaborative Research

RCI resources are vital for complex computations, data management, and collaboration, offering significant benefits across various stages and types of research. These resources, which are *free* (R10), *accessible*, and *open to all* – “anyone can request an account at the RC” (R11), prove invaluable for researchers across disciplines. R05 emphasized their importance even in early research stages, such as proposal writing, where collaboration and data sharing are crucial.

The *ease of access* facilitates data sharing, enabling the exchange of code, data, and analysis results among multiple research teams (R02, R05). Researchers leverage RCI for computational efficiency: GPUs accelerate slow tasks (R12), facilitate essential software for large-scale medical data analysis (R13), and support high-throughput sequencing in biomedical research, advancing disease understanding and drug discovery (R01, R02). RCI supports collaborative projects such as simulating wind effects on bridge engineering, requiring large-scale simulations without physical infrastructure (R06). Additionally, it facilitates adherence to institutional policies and optimizes resource use (R07).

Takeaway: RCIs are indispensable for data sharing and collaboration, meeting the high computational and accessibility needs of interdisciplinary research.

5.3. Ad-hoc and Trust-Based Collaboration

12 out of 14 researcher participants mentioned that access control within their research teams is driven by trust and mutual understanding rather than formal policies. They added that access decisions are often flexible and depend on team members' needs and roles. Trust plays a central role, as R10 emphasized, *"You have to have a lot of trust. That's one of the things I make very clear at the beginning. Ultimately, the team is bigger than the individual."* Data and resources are typically shared freely, with the assumption that team members will act in good faith. As R07 remarked, *"There is enough trust between us that we won't do anything with the data,"* while R14 added, *"It's based on complete trust... if you have access to that particular account, you are not misusing the dataset."*

While formal rules may be absent, access is often managed ad-hoc, relying on PIs' discretion (R04) or team needs (R14), which can sometimes slow down work if the responsible individual is unavailable. As R05 noted:

Oh, [access control decisions are] completely arbitrary. So it's just on a per-need basis. And that changes with time. At some point, I had limited permissions to all my folders and then later I decided to make them accessible to everyone.

This ad-hoc, trust-based approach allows data and resources to be shared freely, though some researchers foresee scalability challenges that may necessitate more formalized policies. The absence of structured access control is viewed as a natural outcome of close collaboration and small team sizes. R04 acknowledged that the lack of automation could lead to bottlenecks as *"one person has to grant access."*

In certain cases, participants noted that access control was not a high priority due to the nature of their research field. R08 highlighted this by contrasting their focus with that of cybersecurity experts:

We're not really sensitive to this kind of thing. What we care most about is the successful implementation of the projects... If we're not violating

any rules and policies in the university or by the funding sponsors.

Takeaway: Access control practices within research teams are often informal, relying on interpersonal trust and ad-hoc decision-making. While this approach may be adequate for small collaborations, researcher participants anticipate that the lack of formal mechanisms can lead to increased complexity and compliance challenges as teams grow.

5.4. Inconsistent Privilege Management

Beyond trust-based, ad-hoc access control practices, privilege management in research teams is occasionally more structured, influenced by factors such as data sensitivity, roles, and resource needs. For example, sensitive data, such as human-related datasets, demands stricter access control (R13). Privileges are tied to individual responsibilities within a project, making access decisions role-based. As R02 explained: *"This student is working on this project, so they will have access to data and computational resources needed for the research tasks."*

Despite some level of structure in these decisions, managing privileges effectively presents several challenges. R07 emphasized the need for granular permissions to reduce the risk of errors:

I only need access to one of the folders within every simulation. It doesn't make sense to give me access to the other 63 directories because I could mistakenly delete everything.

Another major issue mentioned by the participants is that access revocation is more reactive than proactive and usually occurs only when a team member graduates or no longer needs access. As R14 explained: *"Policies remain constant for the lifetime of a project."* R09 pointed out that access removal can even be delayed due to oversight – *"but sometimes my PI forgets."* This lack of proactive access revocation allows former team members to continue using resources after leaving (R04), until their email access is revoked (R10). They added:

I could send an email to our RC center to revoke access, but that requires action from me, which is not automatic, so it's hard to do.

Takeaway: While factors such as data sensitivity and the distinct roles of team members demand more structured privilege management, challenges such as improper and delayed access revocation remain.

6. (RQ3) Critical Stakeholder Challenges

This section outlines the challenges administrators face in securely and systematically managing RCIs, alongside the difficulties researchers encounter in effectively utilizing and sharing resources for scientific collaboration (RQ₃).

6.1. Administrative Challenges with Secure and Systematic RCI Management

In our interviews, system administrators were asked about the challenges they encounter in the daily management of RCI and the obstacles they foresee when integrating more secure, systematic solutions. Their responses highlighted several notable challenges, which we discuss below.

6.1.1. Lack of Automation and Systematic Delegation.

A recurring theme among our participants is the significant administrative burden associated with repetitive tasks in user and privilege management within RCIs. This challenge arises from the absence of automation and a systematic delegation of responsibilities to the users within the ecosystem.

Following our discussion on the *shared responsibility model* (§ 4.3), administrators recognize that empowering researchers with greater control over user and access management could reduce administrative oversight and ensure access decisions align more closely with project needs. However, the lack of a structured delegation framework complicates these tasks further, hindering the efficiency of user and access management processes in RCIs.

Beginning with the user account creation process, automation has the potential to streamline the associated operations significantly as admins “*get a lot of account requests*” (A04). However, challenges arise due to interoperability issues, as RCI admins are often “*dependent on central IT*” for user information, and authentication which is outside the control of RCI management (A03).

A similar challenge concerning the maintenance of ongoing user records was also highlighted. Administrators acknowledge that researchers are better positioned to keep user information up to date; however, the lack of a systematic delegation of these responsibilities makes the process unnecessarily burdensome.

Offloading just a little bit of the user management so that, a PI has to look at the list of students that he’s granting access to and say, someone doesn’t work for me anymore, oh, and I hired two more people. (A10)

Building on this discussion, it was also noted that the lack of a collaborative effort between administrators and users goes beyond the issue of administrative burden and poses potential difficulties in ensuring effective access control in RCIs. As A05 pointed out:

It takes the village to really ensure that... Even if you do have perfect access control, you still could put somebody in the wrong group.

This challenge becomes even more pronounced in the context of user access revocation. The difficulty lies in ensuring that access is properly revoked when users leave a project or the university. As A10 noted:

It’s easy to remember to add people to access controls. It’s really hard to get people to remember to remove people who no longer need access.

Takeaway: *Given the scale of users and the dynamically evolving access needs in RCI, the absence of automated processes and formalized delegation of rights to users make it increasingly challenging to maintain effective access control.*

6.1.2. Infrastructure Complexity and Integration Challenges. When discussing the issues with integrating new access control solutions within RCI, administrators highlighted the complexity and diversity of system components as a primary barrier, rather than the lack of available technologies.

RCIs are built on a combination of disparate systems and applications, each managing access control in its own unique way. For instance, data directories are typically governed by POSIX permissions, while job allocation and execution are managed by tools such as Slurm and PAM. This heterogeneity hinders efforts to establish a unified access control standard across platforms, complicating both management and security. As A05 explained:

To me, the biggest incumbrance here is the applications and systems that do the integrations. And those are hard because they are dispersed and sometimes bespoke, too, right?

Moreover, the foundational components that handle these siloed implementations, such as PAM, POSIX ACLs, and Slurm, are inherently complex. Unix-based PAM modules, for instance, are powerful but notoriously challenging to configure, often leading to unpredictable and error-prone access outcomes that, as A10 described, are “*arcane and complicated*” and can become “*a nightmare.*” Similarly, POSIX ACLs face limited support across HPC resources, and their processing demands pose challenges for high-performance contexts (A05).

Takeaway: *The fragmented system design with highly complex components creates a heterogeneous ecosystem where access control is implemented in a siloed manner across components, increasing the likelihood of errors and security vulnerabilities.*

6.1.3. Ad Hoc Development and Extensibility Issues.

Administrators highlight the ad hoc, piecemeal development of RCI as a critical challenge that undermines consistent and secure system management.

As these infrastructures expand organically, they often accumulate outdated systems and require frequent workarounds or “*hacks on top of hacks*” to maintain functionality (A10). As A08 describes, “*It’s almost like Frankenstein*” as this piecemeal development approach results in infrastructures, where disparate, unplanned components are patched together with limited regard for a cohesive security strategy. Security measures often remain an afterthought, implemented sporadically which makes it difficult to maintain standardized and robust security practices across these platforms. Expanding upon the issues of infrastructure complexity and ad hoc development practices, another challenge

that arises while integrating new security solutions is balancing customizability with generalizability. Given the unique and evolving requirements of advanced scientific computing, administrators often resist one-size-fits-all solutions. These environments prioritize adaptability, seeking to incorporate current technology to serve the specific needs of researchers.

However, this adaptability comes with limitations: systems that work well for one institution often fail to translate to others due to varying institutional requirements, security policies, and administrative capabilities. As A10 noted, even systems considered “*best of breed*” cannot easily be transferred due to institutional security differences and the specialized expertise required for configuration.

Takeaway: *The lack of strategic planning and standardized development practices in RCIs with a security-first approach results in unplanned infrastructures, complicating the adoption and implementation of standardized security practices.*

6.1.4. Human-centric Barriers to Security. Complementing the technical and procedural challenges inherent in RCIs, administrators identified several critical human-centric factors that significantly influence security enhancement.

A significant barrier to enhancing security in RCIs is the resistance to “*changing the mindset*” among long-standing administrators. Many administrators have been managing RCI systems for decades, developing strong preferences for established workflows. These individuals often resist new practices, feeling that “*if everything is working, why change something?*” (A09). This mindset can be pervasive, especially in environments where systems have been functioning reliably under long-held processes.

Additionally, the process of implementing new policies often faces the challenge of “*getting everyone to agree to whatever the rules are*” (A08). Administrators and researchers bring varied perspectives on security policies, with some advocating for stricter measures and others preferring flexibility to support their workflows.

This can lead to “*power struggles*” as individuals resist adjustments that may interfere with their established practices, ultimately slowing down consensus-building and policy implementation. In this context, gaining alignment on new security practices requires not only technical adjustments but also efforts to address deeply ingrained preferences and collaborative tensions.

Further expanding upon this resistance, the reliance on undocumented knowledge within RCI environments poses another important barrier to secure management within RCI. With the lack of accessible and formalized documentation of critical procedural details, newly joined administrators face obstacles in implementing new policies or workflows as much of the information remains “*between [the] ears*” of long-standing staff members (A06). They are typically forced to rely on outdated materials, parts of which often state they “*need to be updated*” (A08). Moreover, complex job policies are often embedded in code rather than

clear guidelines, demanding extensive technical expertise for interpretation. As A10 observed, translating these policies into plain language, such as “*jobs shorter than 20 minutes shall not receive more than X resources*” would streamline management, improve knowledge transfer, and ultimately facilitate smoother adoption of new security measures.

Takeaway: *As outdated practices, delayed decision-making, and insufficient policies hinder the implementation of necessary security improvements, addressing these human-centric factors is essential to complementing technical efforts to enhance the overall security posture of RCI environments.*

6.2. Researcher Experienced RCI Issues

During the interviews, our researcher participants shared various challenges they face when using RCIs for collaborative scientific research. These challenges encompass access control, system configuration, usability, and support, underscoring their impact on collaborative scientific efforts and highlighting the need for improvements in RCIs.

6.2.1. Access Control Issues. Researcher participants frequently reported issues related to insufficient security mechanisms, unintended data exposure, and complexities in managing user access in their experiences with RCI systems.

RCIs Lack Granular Access Control. Participants R06, R07, R09, R12, and R14 noted that many HPC systems default to broad data access permissions, leading to unintended exposure of (sensitive) project data. R07 explained issues with directory permissions:

[RCI admins] asked us to copy datasets we want to share with that temporary collaborator into a directory... we ended up knowing that the guy could access any of the directories within our folder. We definitely don't want any external person to access our data set. But still, we have not figured out any solution.

Downstream Risks of Data Exposure. The lack of compartmentalized permissions exposes researchers to the risk of unintentional or malicious data alteration or deletion. As R12 explained, “*It's very likely that one can delete the data of others... I can access my professor's folder and can have access to the others' folders.*” R13 highlighted the risk of accidental deletion or modification, explaining, “*To prevent that from happening, we backup the data to Amazon and then share the data between our team to do the analysis. This kind of works for us.*”

User Challenges in Managing Access and Permissions. Our interviews revealed that researchers recognize the importance of proper access control and the risks of inadequate management. However, they face significant challenges in implementing and maintaining appropriate access privileges. For example, R01 highlighted difficulties in delegating access, stating, “*When we set up that [partition], we asked*

them [RCI admins] to grant access to certain people... if we want to grant access to others, we have to create a ticket and request it.” This reliance on administrative oversight can delay project timelines and create bottlenecks, hindering seamless collaboration.

R02 expressed concerns about the lack of fine-grained access controls, explaining, “I was told that once there is a partition, everyone in my lab will have access to it... I think if [the university] could make this more specific and let the PI control it, it would be better.” Researchers commonly seek more precise control mechanisms where PIs assign permissions based on team roles and responsibilities. However, placing the responsibility solely on PIs can create challenges, such as oversight or unavailability (§ 5.4).

R04 supported shared management, noting, “If there are partitions that affect a particular lab, someone from that lab should take responsibility.” Sharing such duties among lab members could improve security and efficiency.

Takeaway: Researchers are concerned about the risks of data deletion and unintended exposure due to the lack of granular access control. Manual processes and administrative reliance highlight the need for automated, precise tracking systems.

6.2.2. System-level Configuration Issues. Researchers identified key system-level issues within RCI environments, such as long job queue wait times and frequent system shutdowns, leading to frustration.

Frustrating Queue Systems. A prominent concern was the inefficiency of the queue system used for job submissions, hindering research progress (R14). R03 described the Slurm system: “You give your job, it will wait in queue. Once it’s finished, you give it another job.” R14 expressed frustration with long waits, especially for time-sensitive research: “You’re in the queue for three days... you’re just waiting.” Similarly, R04 said that “whenever I check, it’s always occupied.” R13 added that sharing an account among five team members led to competition for resources, recommending better control over job submissions and CPU usage. Although frustrated, R03 suggested potential solutions to ease the bottleneck: “You can use HTC (high throughput computing) GPU. But if you have longer queue, I think we might need a bit more computation power.”

Maintenance and Downtime. Frequent system maintenance and unexpected shutdowns were significant frustrations. R06 shared, “There’s a lot of maintenance... sometimes it says when there’s going to be an outage, but it doesn’t always say when there’s maintenance... that’s two days where I can’t really make progress.”

In-house Computing and Private Partitions. To address frequent RCI disruptions, some researchers turned to in-house or private computing solutions for greater control and flexibility. R04 explained that the university’s clusters often lead to underutilized GPU capacity, with “a lot of GPU is getting wasted.” By managing internal servers, R04’s team aim to improve GPU allocation and hardware use.

Similarly, R09 faced software compatibility issues with shared resources (inability to install specific Python packages on the university server), prompting them to rely on local machines. In contrast, R10 opted for a private partition on the university’s infrastructure, balancing cost and convenience with priority access and dedicated support. As they noted, “if I want to submit a job... other jobs [are] kicked out.” This provides a middle ground between centralized and self-managed systems, despite the added costs.

Takeaway: Reliance on shared resources causes competition, delays, and frustration, prompting researchers to seek alternative solutions such as in-house computing or private partitions for greater control and reduced dependence on centralized systems.

6.2.3. Usability and Support Issues. Several researchers highlighted usability and support challenges that hindered their effective use of RCIs, including inadequate training and orientation, and inconsistent support from technical staff.

Training Barriers. R06 expressed frustration with inadequate training, noting “the only useful thing was how to log in.” They struggled with tasks such as creating job scripts, relying on colleagues for help – “I kind of had to ask other people in my lab group and kind of stumble around.” R08’s student avoided HPC due to the need for extra training and scheduling issues. R09 mentioned feeling unprepared in their new lead researcher role, stating, “there is no really good orientation to how to manage the different resources.”

Support Gaps. R05 mentioned that while RCI staff were generally helpful, their high workloads led to delayed responses: “Sometimes they just drop the ball and don’t reply for weeks.” R01 and R14 reported ongoing issues with access permissions, with admins failing to resolve repeated requests. R10 summarized the overall disappointment toward available RCI support when they said:

We wish we had some more help... but of course, usually these research computing centers don’t have the resources to really help us in a meaningful way. So we are always forced to figure things more or less out by ourselves.

Takeaway: Insufficient training and delayed support from RCI staff hinder effective use of systems, forcing researchers to rely on peers or self-solving, which leads to inefficiencies and frustration in research workflows.

6.2.4. Sentiment Analysis. While researchers discussed various challenges associated with their current RCI systems experiences, only 5 out of 14 participants expressed negative sentiments about their overall satisfaction. For instance, R03 and R14 expressed frustration with job submission queues, while R05, though frustrated, put it more mildly, saying:

Enough, enough is the word that is doing a lot of heavy lifting here. It could be worse.

R06 noted they had resolved small issues but still found things manageable, saying, “I will keep using it... because at this point, I figured most of those things out too.” R09 took a pragmatic approach, stating, “If I end up having to just remote run it locally on my machine, I’ll do that.”

On the positive side, R01 stated, “Overall, it’s pretty positive. I think the resources are abundant... and of course, sometimes you have to be queued for a while, but overall, it’s still pretty decent.” Similar to R01, R13 emphasized the responsiveness of the RCI team, sharing, “We’re quite satisfied... they respond really fast and can resolve all the problems you have very quickly.” R04 expressed, *I feel like [RCI] has done a great job with whatever resources we have.* and R11 shared, “They have everything I need.”

Takeaway: *Despite various challenges, many researchers recognize the value and utility of the RCIs, reflecting a complex balance between limitations and continued reliance on these systems.*

7. Discussion

In this section, we critically analyze our findings and present key recommendations on RCI access control and system interaction shaped by the experiences, challenges, and needs of our participants. These recommendations aim to guide future research in improving both security and usability in RCIs, while fostering a more efficient and collaborative scientific environment.

7.1. Access Control Requirements in RCIs

Here, we outline key access control requirements that provide actionable guidance and serve as a foundation for secure and effective privilege management within RCIs.

(R_1) **Formal Modeling of Ownership:** Resource ownership in RCIs exhibits a *multi-layered structure* shaped by diverse resource types, distinct user roles and responsibilities, and unique collaborative needs (§ 4.2). Unlike traditional environments such as cloud computing or enterprise systems, which primarily rely on either admin-driven (e.g., RBAC [58]) or owner-driven (e.g., DAC [21]) access control approaches, RCIs support intricate relationships among users, administrators, and resources. These ecosystems must manage resources across stakeholders ranging from individual researchers to group-based projects and institution-wide administrators, each with a distinct scope of resource ownership and access needs. Therefore, it is critical to formally model the multi-layered ownership to enable policy-driven granular access control and discretionary resource sharing.

(R_2) **Multi-layered Delegation and Conflict Resolution:** The complex ecosystem of RCIs requires a structured yet flexible *shared responsibility model* (§ 6.1.1) for privilege management, as researchers expressed frustration with the

lack of nuanced permission settings, emphasizing the need for fine-grained control over access to resources (§ 6.2.1). Such a model should align with the *multi-layered ownership* framework (R_1) and support dynamic delegation according to academic hierarchy and role precedence (§ 5.1), enabling each stakeholder to uniquely contribute based on their distinct roles and decision-making scopes.

Administrators can adapt rule-based delegation frameworks (e.g., [59], [60]) to define these policies to ensure appropriate distribution of access privileges across individual, group, and institutional levels, mitigating excessive privilege allocation. At the same time, mechanisms for real-time conflict detection and resolution (e.g., [61], [62], [63]) must be incorporated to ensure that resource ownership and priority—delegated by administrators to PIs, and from PIs to collaborators—remains secure and aligned with overarching institutional, legal, and sponsored-issued policies while meeting project-specific needs.

(R_3) **Automated Context-aware Privilege Revocation:** As RCIs support a large user base with evolving access needs, automated revocation mechanisms are crucial for ensuring timely and secure removal of access as roles and project affiliations change (§ 5.4, § 6.1.1). While access revocation has been extensively studied in the literature (e.g., [64], [65], [66]), enforcing it in highly collaborative environments such as RCIs requires balancing security with productivity. Efficient revocation must minimize administrative overhead while ensuring uninterrupted resource availability through timely access renewals.

A promising approach is to contextualize privilege management by defining access within RCI-specific temporal units—such as collaborations, projects, roles, or tasks within projects—while considering their temporal and contextual interdependencies. Establishing clear start and end points for privileges allows for automated revocation. Additionally, optimizing access renewal can be achieved by adjusting the granularity of the privilege context—whether tied to something as small as a task, or a long-standing project.

(R_4) **Unified Privilege Management:** A cohesive approach to privilege management is essential to address the fragmented and siloed access control practices prevalent across complex, heterogeneous components (§ 6.1.2, § 6.1.3). Support for unique management policies across diverse resources can be achieved through an abstraction layer over the underlying infrastructure. However, unifying privilege management through abstraction presents challenges due to varying institutional security policies, compliance requirements, and administrative preferences. To ensure adoption, this abstraction must be designed with administrative flexibility in mind, allowing modular customization to align with different institutional needs. While federated access management has been explored (e.g., [67], [68]), its application to RCIs requires further investigation, particularly in addressing scalable adaptation and policy conflict resolution across multiple institutions.

7.2. Recommendations on improving Human-System Interaction in RCIs

In this section, we go beyond access control and provide key recommendations—derived from researcher-shared challenges (§ 6.2)—on enhanced usability and system interactions within RCIs.

System-interaction Enhancements. Long wait times for critical resources, such as GPUs, were a significant pain point, where delays often negatively impacted project timelines (§ 6.2.2). To reduce scheduling delays, RCIs should explore hybrid queue systems or integrate high-throughput computing (HTC) solutions (R03), such as HTCCondor [69]. However, adopting HTC solutions to maximize resource availability must balance with the computational efficiency offered by HPC systems.

Furthermore, enhanced resource monitoring—leveraging real-time visualization tools—could improve resource availability management, ensuring better alignment with the evolving demands of research projects. While admins utilize tools like Open XDMoD [70], there is a need for monitoring solutions tailored for non-admin researchers. Providing researchers with intuitive tools would empower them to track resource availability more effectively, enabling proactive adjustments based on the specific computational needs of individuals, groups and projects.

Open-source tools such as ColdFront [71] are gaining popularity for their usability, allowing PIs to create projects, request allocations for RCI resources, and perform periodic reviews to ensure proper revocation of privileges. While such tools improve accessibility, to ensure security, future efforts must explore how to integrate more nuanced access control requirements (§ 7.1) within these tools to balance usability with security.

Additionally, participants (R02, R06) recommended introducing transparent communication protocols for planned downtimes and system updates to minimize disruptions and ensure continuity of research. Establishing user advocacy channels could enhance communication and ensure that user concerns are heard and addressed.

Usability Enhancements. Participants highlighted challenges from inadequate training (§ 6.2.3), including R06’s difficulty with job scripts, R08’s avoidance of HPC, and R09’s unpreparedness in a lead researcher role. Role-specific, scenario-based training could simplify the adoption of complex systems. PEARC’s [72] session on workforce development, training, diversity, and education underscores the importance of a community-driven approach that includes academia, government, and industry.

8. Conclusion

In this paper, we presented the first in-depth qualitative study on scientific collaboration and access control practices within Research Computing Infrastructures (RCIs). Through comprehensive, semi-structured interviews with

24 key RCI stakeholders—comprising PIs, graduate researchers, and system administrators—we provide multi-perspective insights into the current practices, challenges, and requirements for secure and systematic resource sharing and access management in scientific collaborations.

Our findings reveal the operational complexities of user, resource, and privilege management as practiced by system administrators, alongside the project-specific, trust-driven access control dynamics that shape collaborative research teams. Overall, this work seeks to uncover previously unexplored security, usability, and access control challenges within RCIs, while capturing participant perspectives to inform future advancements. By providing a valuable framework for improving access control and broader research computing practices, this study contributes to sustaining and strengthening the overall security posture of RCIs.

Acknowledgment

We thank the anonymous reviewers and our shepherd for their thoughtful and constructive feedback. We are especially grateful to the participants for generously sharing their time and insights into the complexities of scientific collaboration and privilege management.

This work was partly supported by the National Science Foundation (NSF) under grant NSF-CICI-2232911 and the Institute of Information & Communications Technology Planning & Evaluation (IITP) through the following grants: RS-2024-004398199 and RS-2024-00442085.

References

- [1] R. Kitchin, “Big data, new epistemologies and paradigm shifts,” *Big data & society*, vol. 1, no. 1, p. 2053951714528481, 2014.
- [2] M. Parashar, “Democratizing science through advanced cyberinfrastructure,” *Computer*, vol. 55, no. 9, pp. 79–84, 2022.
- [3] T. J. Boerner, S. Deems, T. R. Furlani, S. L. Knuth, and J. Towns, “Access: Advancing innovation: Nsf’s advanced cyberinfrastructure coordination ecosystem: Services & support,” in *Practice and Experience in Advanced Research Computing*, pp. 173–176, 2023.
- [4] E. Strubell, A. Ganesh, and A. McCallum, “Energy and policy considerations for modern deep learning research,” in *Proceedings of the AAAI Conf. on AI*, pp. 13693–13696, 2020.
- [5] P. Muir, S. Li, S. Lou, D. Wang, D. J. Spakowicz, L. Salichos, J. Zhang, G. M. Weinstock, F. Isaacs, J. Rozowsky, *et al.*, “The real cost of sequencing: scaling computation to keep pace with data generation,” *Genome biology*, vol. 17, pp. 1–9, 2016.
- [6] Z. D. Stephens, S. Y. Lee, F. Faghri, R. H. Campbell, C. Zhai, M. J. Efron, R. Iyer, M. C. Schatz, S. Sinha, and G. E. Robinson, “Big data: astronomical or genomics?,” *PLoS biology*, vol. 13, no. 7, 2015.
- [7] L. D. Stein, “Towards a cyberinfrastructure for the biological sciences: progress, visions and challenges,” *Nature Reviews Genetics*, vol. 9, no. 9, pp. 678–688, 2008.
- [8] “TOP500 List.” <https://www.top500.org/>. (Accessed on 11/31/2024).
- [9] “Bu and federal investigation underway into hacking of framingham heart study data.” <https://www.bu.edu/articles/2024/investigation-into-hacking-of-fhs-data/>. Accessed: 2025-03-18.
- [10] J. Reed, “White house mandates stricter cybersecurity for R&D institutions.” <https://securityintelligence.com/news/white-house-mandates-stricter-cybersecurity-for-research-and-development/>, 2024. Accessed: 2024-11-03.

- [11] "Cybersecurity innovation for cyberinfrastructure (cici)." <https://new.nsf.gov/funding/opportunities/cybersecurity-innovation-cyberinfrastructure-cici>, 2024. Accessed: 2024-11-03.
- [12] A. Jezghani, S. Sarajlic, M. Brandon, N. Bright, M. Belgin, G. Beyer, C. Blanton, P. Buffington, J. E. Coulter, R. Lara, *et al.*, "Phoenix: The revival of research computing and the launch of the new cost model at georgia tech," in *Practice and Experience in Advanced Research Computing*, pp. 1–9, 2022.
- [13] D. M. Jennewein, J. Lee, C. Kurtz, W. Dizon, I. Shaeffer, A. Chapman, A. Chiquete, J. Burks, A. Carlson, N. Mason, *et al.*, "The sol supercomputer at arizona state university," in *Practice and Experience in Advanced Research Computing*, pp. 296–301, 2023.
- [14] "Research computing infrastructure (rci) — ohio state college of medicine." <https://medicine.osu.edu/research/research-information-technology/services-and-products/research-computing-infrastructure>. Accessed: 2024-09-02.
- [15] "Research computing infrastructure — princeton neuroscience institute." <https://pni.princeton.edu/research-areas/research-computing-infrastructure>. Accessed: 2024-09-02.
- [16] "About rc core facilities — arizona state university." <https://cores.research.asu.edu/research-computing/about>. Accessed: 2024-09-02.
- [17] "Research computing infrastructure - center for computational research - university at buffalo." <https://www.buffalo.edu/ccr/services/high-performance-computing.html>. Accessed: 2024-09-02.
- [18] GIGALIGHT, "Why infiniband networks are crucial in high-performance computing data centers?." <https://gigalight.medium.com/why-infiniband-networks-are-crucial-in-high-performance-computing-data-centers-cb74d37516c8>. Accessed: 2024-09-02.
- [19] A. B. Yoo, M. A. Jette, and M. Grondona, "Slurm: Simple linux utility for resource management," in *Workshop on job scheduling strategies for parallel processing*, pp. 44–60, Springer, 2003.
- [20] V. Samar, "Unified login with pluggable authentication modules (pam)," in *Proceedings of the 3rd ACM conference on Computer and communications security*, pp. 1–10, 1996.
- [21] Stack Exchange Information Security, "Why is linux filesystem considered dac and not mac?" <https://security.stackexchange.com/questions/199153/why-is-linux-filesystem-considered-dac-and-not-mac>, 2020. Accessed: 2024-11-13.
- [22] G. Burleson, M. Machado, and I. Aranda, "'engineering for global development' in academic institutions: An initial review of learning opportunities across four global regions," in *2021 World Engineering Education Forum/Global Engineering Deans Council (WEEF/GEDC)*, pp. 153–158, IEEE, 2021.
- [23] N. L. Cole, S. Reichmann, and T. Ross-Hellauer, "Toward equitable open research: stakeholder co-created recommendations for research institutions, funders and researchers," *Royal Society Open Science*, vol. 10, no. 2, p. 221460, 2023.
- [24] K. Chapman, G. Ruan, E. Tuna, A. Walsh, and E. Wernert, "Scholarly data share: A model for sharing big data in academic research," in *PEARC*, pp. 1–8, 2022.
- [25] K. Chapman, G. Ruan, E. Tuna, A. Walsh, and E. A. Wernert, "Scholarly data share 2.0: Granular access to research data," in *PEARC*, pp. 177–180, 2023.
- [26] M. Christie, S. Marru, S. Pamidighantam, I. Ranawaka, and D. Wanipurage, "Airavata data catalog: A multi-tenant metadata service for efficient data discovery and access control," in *Practice and Experience in Advanced Research Computing*, pp. 181–185, 2023.
- [27] M. J. H. Faruk, J. Basney, and J. Q. Cheng, "Blockchain-based decentralized verifiable credentials: Leveraging smart contracts for privacy-preserving authentication mechanisms to enhance data security in scientific data access," in *2023 IEEE International Conference on Big Data (BigData)*, pp. 5493–5502, IEEE, 2023.
- [28] S. Sivagnanam, S. Yeu, K. Lin, S. Sakai, F. Garzon, K. Yoshimoto, K. Prantzas, D. Upadhyaya, A. Majumdar, S. S. Sahoo, *et al.*, "Towards building a trustworthy pipeline integrating neuroscience gateway and open science chain," *Database*, 2024.
- [29] J. Anderson and K. Keahey, "Migrating towards single sign-on and federated identity," in *Practice and Experience in Advanced Research Computing*, pp. 1–8, 2022.
- [30] B. Aydemir, J. Basney, B. Bockelman, J. Gaynor, and D. Weitzel, "Sciauth: A lightweight end-to-end capability-based authorization environment for scientific computing," in *Practice and Experience in Advanced Research Computing*, pp. 1–5, 2022.
- [31] A. Withers, B. Bockelman, D. Weitzel, D. Brown, J. Gaynor, J. Basney, T. Tannenbaum, and Z. Miller, "Scitokens: capability-based secure access to remote scientific data," in *Proceedings of the practice and experience on advanced research computing*, pp. 1–8, 2018.
- [32] K. Kuznia, D. Shah, G. Speyer, and J. Yalim, "Artificial intelligence to classify and detect masquerading users on hpc systems from shell histories," in *Practice and Experience in Advanced Research Computing*, 2022.
- [33] Y. Wu, W. K. Edwards, and S. Das, "Sok: Social cybersecurity," in *2022 IEEE Symposium on Security and Privacy (SP)*, pp. 1863–1879, IEEE, 2022.
- [34] S. Das, T. H.-J. Kim, L. A. Dabbish, and J. I. Hong, "The effect of social influence on security sensitivity," in *10th Symposium On Usable Privacy and Security (SOUPS 2014)*, pp. 143–157, 2014.
- [35] F. B. Kokulu, A. Soneji, T. Bao, Y. Shoshitaishvili, Z. Zhao, A. Doupé, and G.-J. Ahn, "Matched and mismatched socs: A qualitative study on security operations center issues," in *Proceedings of the 2019 ACM SIGSAC conference on computer and communications security*, pp. 1955–1970, 2019.
- [36] I. Ion, R. Reeder, and S. Consolvo, "{... No} one can hack my {Mind}": Comparing expert and {Non-Expert} security practices," in *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*, pp. 327–346, 2015.
- [37] M. Fagan and M. M. H. Khan, "Why do they do what they do?: A study of what motivates users to (not) follow computer security advice," in *Twelfth symposium on usable privacy and security (SOUPS 2016)*, pp. 59–75, 2016.
- [38] B. L. BERG, "Qualitative research methods for the social sciences," 2001.
- [39] R. Werlinger, K. Hawkey, K. Muldner, P. Jaferian, and K. Beznosov, "The challenges of using an intrusion detection system: is it worth the effort?," in *Proceedings of the 4th symposium on Usable privacy and security*, pp. 107–118, 2008.
- [40] P. Jaferian, H. Rashtian, and K. Beznosov, "To authorize or not authorize: helping users review access policies in organizations," in *10th Symposium On Usable Privacy and Security (SOUPS 2014)*, pp. 301–320, 2014.
- [41] A. Soneji, F. B. Kokulu, C. Rubio-Medrano, T. Bao, R. Wang, Y. Shoshitaishvili, and A. Doupé, "'flawed, but like democracy we don't have a better system': The experts' insights on the peer review process of evaluating security papers," in *2022 IEEE Symposium on Security and Privacy (SP)*, pp. 1845–1862, IEEE, 2022.
- [42] D. Votipka, R. Stevens, E. Redmiles, J. Hu, and M. Mazurek, "Hackers vs. testers: A comparison of software vulnerability discovery processes," in *2018 IEEE Symposium on Security and Privacy (SP)*, pp. 374–391, IEEE, 2018.
- [43] C. P. Lee, P. Dourish, and G. Mark, "The human infrastructure of cyberinfrastructure," in *Proceedings of the 20th conference on Computer supported cooperative work*, pp. 483–492, 2006.
- [44] M. J. Bietz, E. P. Baumer, and C. P. Lee, "Synergizing in cyber-infrastructure development," *Computer Supported Cooperative Work (CSCW)*, vol. 19, pp. 245–281, 2010.

- [45] P. Chityala, C. Costa, J. Wernert, and C. Stewart, "Cyberinfrastructure value: A survey on perceived importance and usage," in *Practice and Experience in Advanced Research Computing*, pp. 1–4, 2022.
- [46] B. Yeager, J. Yalim, S. Knuth, A. Romanella, C. Reidy, and B. Nickell, "Identifying high-performance computing needs at member institutions in a regional consortium," in *Practice and Experience in Advanced Research Computing 2024: Human Powered Computing*, pp. 1–5, 2024.
- [47] R. Atkinson and J. Flint, "Accessing hidden and hard-to-reach populations: Snowball research strategies," *Social research update*, vol. 33, no. 1, pp. 1–4, 2001.
- [48] C. Parker, S. Scott, and A. Geddes, "Snowball sampling," *SAGE research methods foundations*, 2019.
- [49] "Replication package: Investigating the complexities of scientific collaboration and privilege management within research computing infrastructures," <https://github.com/sefcom/Frankenstein/tree/master>. Accessed: 2025-03-13.
- [50] S. Rädiker and U. Kuckartz, "Focused analysis of qualitative interviews with maxqda," *MAXQDA Press*, 2020.
- [51] V. Braun and V. Clarke, *Thematic analysis*. American Psychological Association, 2012.
- [52] R. Hodson, *Analyzing documentary accounts*. No. 128, Sage, 1999.
- [53] M. L. McHugh, "Interrater reliability: the kappa statistic," *Biochemia medica*, vol. 22, no. 3, pp. 276–282, 2012.
- [54] A. S. University, "Human subjects review process — research compliance," <https://researchcompliance.asu.edu/human-subjects/review-process/>. Accessed: 2025-03-13.
- [55] J. Pater, A. Coupe, R. Pfafman, C. Phelan, T. Toscos, and M. Jacobs, "Standardizing reporting of participant compensation in hci: A systematic literature review and recommendations for the field," in *Proceedings of the 2021 CHI conference on human factors in computing systems*, pp. 1–16, 2021.
- [56] J. Mink, H. Kaur, J. Schmöser, S. Fahl, and Y. Acar, "{Security} is not my field, {I'm} a stats {guy}": A qualitative root cause analysis of barriers to adversarial machine learning defenses in industry," in *32nd USENIX Security Symposium*, pp. 3763–3780, 2023.
- [57] M. Hausenblas, *Learning Modern Linux*. O'Reilly Media, Inc., 2022.
- [58] R. S. Sandhu, "Role-based access control," in *Advances in computers*, vol. 46, pp. 237–286, Elsevier, 1998.
- [59] L. Zhang, G.-J. Ahn, and B.-T. Chu, "A rule-based framework for role-based delegation and revocation," *ACM Transactions on Information and System Security*, vol. 6, no. 3, pp. 404–441, 2003.
- [60] E. Barka, R. Sandhu, *et al.*, "A role-based delegation model and some extensions," in *Proceedings of the 23rd National Information Systems Security Conference*, vol. 4, pp. 49–58, NIST Baltimore, 2000.
- [61] S. Benferhat, R. El Baida, and F. Cuppens, "A stratification-based approach for handling conflicts in access control," in *Proceedings of the eighth ACM symposium on Access control models and technologies*, pp. 189–195, 2003.
- [62] T. Jaeger, R. Sailer, and X. Zhang, "Resolving constraint conflicts," in *Proceedings of the ninth ACM symposium on Access control models and technologies*, pp. 105–114, 2004.
- [63] M. Koch, L. V. Mancini, and F. Parisi-Presicce, "Conflict detection and resolution in access control policy specifications," in *International Conference on Foundations of Software Science and Computation Structures*, pp. 223–238, Springer, 2002.
- [64] S. Jahid, P. Mittal, and N. Borisov, "Easier: Encryption-based access control in social networks with efficient revocation," in *Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security*, pp. 411–415, 2011.
- [65] G. Yu, X. Zha, X. Wang, W. Ni, K. Yu, P. Yu, J. A. Zhang, R. P. Liu, and Y. J. Guo, "Enabling attribute revocation for fine-grained access control in blockchain-iot systems," *IEEE Transactions on Engineering Management*, vol. 67, no. 4, pp. 1213–1230, 2020.
- [66] V. D. Gligor, "Review and revocation of access privileges distributed through capabilities," *IEEE Transactions on Software Engineering*, no. 6, pp. 575–586, 1979.
- [67] C. E. Rubio-Medrano, Z. Zhao, A. Doupé, and G.-J. Ahn, "Federated access management for collaborative network environments: Framework and case study," in *Proceedings of the 20th ACM Symposium on Access Control Models and Technologies*, pp. 125–134, 2015.
- [68] P. Fremantle, B. Aziz, J. Kopecký, and P. Scott, "Federated identity and access management for the internet of things," in *2014 International Workshop on Secure Internet of Things*, pp. 10–17, 2014.
- [69] "Htcondor software suite," <https://htcondor.org/>. Accessed: 2025-03-24.
- [70] J. T. Palmer, S. M. Gallo, T. R. Furlani, M. D. Jones, R. L. DeLeon, J. P. White, N. Simakov, A. K. Patra, J. Sperhac, T. Yearke, *et al.*, "Open xdm: A tool for the comprehensive management of high-performance computing resources," *Computing in Science & Engineering*, vol. 17, no. 4, pp. 52–62, 2015.
- [71] A. Bruno and D. Sajdak, "Coldfront: Resource allocation management system," in *Practice and Experience in Advanced Research Computing 2021: Evolution Across All Dimensions*, pp. 1–5, 2021.
- [72] PEARC, "Practice and experience in advanced research computing," <https://pearc.acm.org/pearc24/>, 2024. Accessed: 2024-11-13.

Appendix A. Questionnaire for Researchers

The questionnaire tailored to the Researchers was divided into four distinct sections as follows. The overlapping questions with administrators have been marked as $[R + A]$.

Understanding Scientific Collaboration

- 1) $[R + A]$ Can you describe the notion of a "research team" from a research collaboration perspective, specifically *why researchers form teams and how these teams are formed*?
- 2) From the team management perspective, is there any specific structure of a team?
- 3) Are you part of such a team/teams?
 - a) What kind of collaborative tasks do you engage in for scientific research?

Understanding Resource Sharing

- 4) Do you share any digital resources with the other research team members, including both computing and non-computing resources?
 - a) *[Follow up] (If they mention computing resources)*
 - i) What computing resources do you share?
 - ii) Is it shared only internally within your institution, or externally as well? Why?
 - b) *[Follow up] (Otherwise)*
 - i) Are there any specific factors that have been preventing you from sharing computing resources?
 - ii) If given the facility, what computing resources would you like to share?
- 5) On a high level, how do you manage access permissions to digital resources within your collaborative team?
- 6) Do you keep track of what resources are being shared by the different members of the research team?

- a) How do you keep track of this? Any software/tool in this context?
 - b) Is the tracking process “manual” or “automated”? Can you elaborate on the exact process?
 - c) Is there any dedicated person on the team for this?
 - d) How frequently are the sharing records updated?
 - e) If no tracking, then why?
- 7) In your collaborative teams, how is it decided who is going to access what resources? Are there *any specific rules or policies* in place?
- a) [Follow up] (If answered yes) Can you elaborate more on the current rules?
 - b) [Follow up] (If answered no) Why do you think you do not have a rule system in place? Is there any benefit of implementing rules?
- 8) How are these constraints or rules usually formed or should be formed? Is it the *discretion of the PI*, or a *consensus of the team*, or *something else*?
- 9) How do you keep track of whether these access rules are being followed?
- 10) What do you think usually causes such changes in the rules? Is there any systematic approach for such changes?

Understanding Challenges and Needs

- 11) [R + A] Are there any factors that contribute to your feelings of frustration when interacting with the RC environment in your day-to-day life?
- 12) [R + A] Could there be any better support for more secure collaboration in RC?
- 13) [R + A] In the context of access management, how much of the administrative privileges or autonomy can be delegated to the users?
- 14) Do you think the RC infrastructure should support such features through which you can implement/enforce custom rules for your privately owned resources?
- 15) [R + A] Are there any aspects that should be enhanced/updated to improve the overall user experience of the RC infrastructure?

Understanding Satisfaction and Future Plan

- 16) [R + A] Are you satisfied enough with the RC infrastructure that you will continue using it for the foreseeable future?

Appendix B.

Questionnaire for Administrators

The questionnaire tailored to the Administrators was divided into four distinct sections as follows.

Understanding Existing Infrastructure

- 1) What digital resources are currently being managed by the Research Computing infrastructure?
- 2) What are the processes of how these different resources are managed? Can you tell us about the different software/tools you use in such a context?
- 3) How are these digital resources restricted on an individual user level? Can all users access all the same resources in the system or otherwise?

Understanding Support for Collaboration

- 4) How many different types of user accounts are there in the RC System?
 - a) [Follow up] Do the sponsor accounts get any special privileges on restricting access to the accounts they are sponsoring?
- 5) How can each user in the system be characterized?
- 6) What is the user management process in the system, including the tools/software involved?
- 7) [R + A] Can you describe the notion of a “team” from a research collaboration perspective?
- 8) From a system administrator’s perspective, how are these teams usually managed?
- 9) In the current infrastructure, is there any way that users can share resources among themselves?
 - a) [Follow up] (If yes) Can you elaborate on what resources can be shared, with whom, and what is the process of sharing?
 - b) [Follow up] (If no) Can you think of any ways the resources could possibly be shared?
- 10) Does the current system support the users proposing/implementing their own rules/constraints on how resources can be shared between team members?
 - a) [Follow up] (If yes) How do they achieve it? Are there any such rules?
 - b) [Follow up] (If no) What could be the possible benefit of this feature from an administrative perspective? Can you think of any possible issues?

Understanding Challenges and Needs

- 11) [R + A] Are there any factors that contribute to your feelings of frustration while administering the RC environment in your day-to-day life?
- 12) [R + A] What kinds of support for collaboration would be more effective in research computing infrastructures?
- 13) [R + A] In the context of access management, how much administrative privileges or autonomy can be delegated to the users?
- 14) Are there any features that are not implemented yet, but could potentially ease the process of managing resources and giving users access to those resources?
 - a) [Follow up] Can you think of any already existing tools or software in this context?
 - a) If a dedicated access control system is to be implemented for the research computing infrastructure, what features should it have?
- 15) [R + A] Are there any aspects of the administrative policies that should be enhanced/updated to improve the overall experience of the users of the RC infrastructure?

Understanding Satisfaction and Future Plan

- 16) [R + A] Can you tell us about the future plans of the RC Infrastructure that you are aware of to enable better service to its potential users?

Appendix C.

Meta-Review

The following meta-review was prepared by the program committee for the 2025 IEEE Symposium on Security and Privacy (S&P) as part of the review process as detailed in the call for papers.

C.1. Summary

This paper presents a qualitative study on the collaboration, challenges, needs and security concerns of researchers and system administrators in Research Computing Infrastructures (RCIs). The study identifies key challenges that include trust-based access control, lack of automated processes and fragmented system designs. The paper provides recommendations to improve security, usability, and access control requirements in RCIs.

C.2. Scientific Contribution

- Provides a Valuable Step Forward in an Established Field
- Establishes a New Research Direction

C.3. Reasons for Acceptance

- 1) The study documents the human and administrative challenges of access control in RCI and highlights the need for usable solutions that address the difficulties faced by both researchers and administrators.
- 2) The paper presents key findings, including that trust-based access control and the decentralized nature of RCI management may conflict with IRB policies when sharing datasets containing sensitive information.