When the Weakest Link is Strong: Secure Collaboration in the Case of the Panama Papers

Susan E. McGregor, ***Columbia Journalism School;*** **Elizabeth Anne Watkins,** ***Columbia University;*** **Mahdi Nasrullah Al-Ameen and Kelly Caine,** ***Clemson University;*** **Franziska Roesner,** ***University of Washington***

https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/mcgregor

This paper is included in the Proceedings of the 26th USENIX Security Symposium

August 16–18, 2017 • Vancouver, BC, Canada

ISBN 978-1-931971-40-9

Open access to the Proceedings of the 26th USENIX Security Symposium is sponsored by USENIX

When the Weakest Link is Strong: Secure Collaboration in the Case of the Panama Papers

Susan E. McGregor
Columbia Journalism School

Elizabeth Anne Watkins
Columbia University

Mahdi Nasrullah Al-Ameen
Clemson University

Kelly Caine
Clemson University

Franziska Roesner
University of Washington

Abstract

Success stories in usable security are rare. In this paper, however, we examine one notable security success: the year-long collaborative investigation of more than two terabytes of leaked documents during the "Panama Papers" project. During this effort, a large, diverse group of globally-distributed journalists met and maintained critical security goals–including protecting the source of the leaked documents and preserving the secrecy of the project until the desired launch date–all while hundreds of journalists collaborated remotely on a near-daily basis.

Through survey data from 118 participating journalists, as well as in-depth, semi-structured interviews with the designers and implementers of the systems underpinning the collaboration, we investigate the factors that supported this effort. We find that the tools developed for the project were both highly useful and highly usable, motivating journalists to use the secure communication platforms provided instead of seeking workarounds. We also found that, despite having little prior computer security experience, journalists adopted—and even appreciated—the strict security requirements imposed by the project leads. We also find that a shared sense of community and responsibility contributed to participants' motivation to meet and maintain security requirements. From these and other findings, we distill lessons for socio-technical systems with strong security requirements and identify opportunities for future work.

1 Introduction

On April 3, 2016, a coordinated network of dozens of news organizations around the world [32] began publishing stories based on a set of year-long investigations into the uses of offshore funds by clients of the Panamanian law firm Mossack Fonseca. The revelations contained in these "Panama Papers" led to the ouster of Icelandic Prime Minister Sigmundur David Gunnlaugsson [17], and helped instigate investigations from Argentina and Australia to Canada, Denmark, France, India, Indonesia, Mexico, Pakistan, and others [42].

Facilitated by the International Consortium of Investigative Journalists (ICIJ), the Panama Papers project [31] represents a uniquely positive security case study, wherein systems designed, implemented, and managed by a handful of ICIJ staffers helped meet and maintain the organization's security goals for the project. While it is impossible to state definitively that this (or any) system could *not* have been compromised, ICIJ's efforts appear to have been successful in maintaining their primary security goals, including: (1) protecting the identity of the source of the Panama Papers' documents (2) maintaining control of the documents within their network of collaborators and preventing their early public disclosure, (3) protecting the documents themselves from attackers (e.g., the companies, criminals and political figures they implicated), and, finally (4) keeping the investigation itself a secret for over a year. Remarkably, all of this was achieved while supporting the collaborative analysis of the documents by nearly 400 journalist-contributors worldwide, who communicated regularly across time zones and language barriers.

In the computer security literature and beyond, users are often referred to as "the weakest link" in security systems (e.g., [26, 48, 50]). Recent case studies on activist organizations and NGOs [21, 39, 43], for example, highlight such security failures in context. Through examination of the Panama Papers project, then, we seek to learn (1) what technical and human factors facilitated the successful preservation of the project's security goals and, (2) what lessons can be drawn from this case study to support the development of similarly effective processes for both journalistic collaborations and secure, usable systems in general. For while the technical systems used in the Panama Papers project did not necessarily incorporate all technical security best practices, our investigation helps illuminate how the systems' hundreds of users were nevertheless able to collaborate securely over a long period of time.

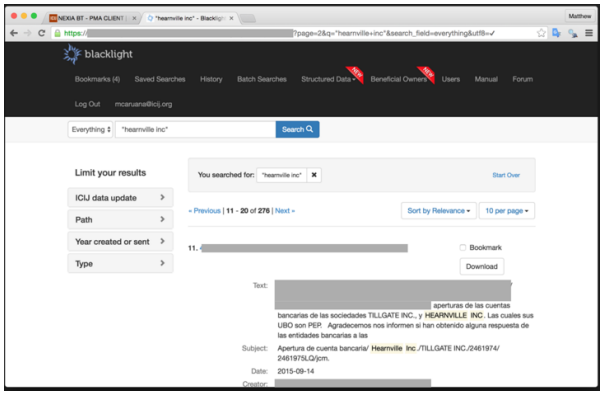

Figure 1: **Blacklight.** Screenshot of the document search platform. *Courtesy: ICIJ.*

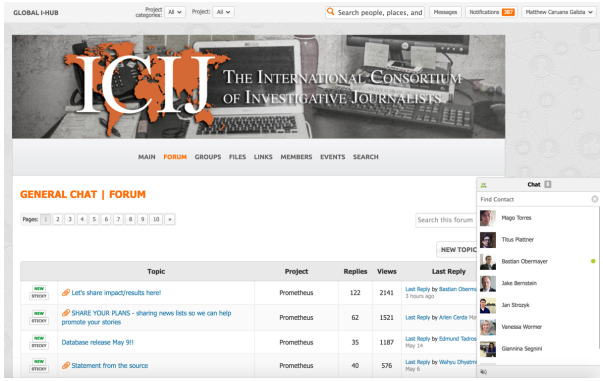

Figure 2: **I-Hub.** Screenshot of the collaboration and communication platform. *Courtesy: ICIJ.*

To uncover the factors that contributed to the Panama Papers' security success, we (1) analyze survey data collected from 118 journalists involved in the project, and (2) conduct in-depth, semi-structured interviews with the designers and implementers of the technical systems and collaborative processes underpinning the Panama Papers collaboration. These systems[1] included:

- **Blacklight** (Figure 1), a document-search platform where contributing journalists could access the leaked documents.
- **I-Hub** (Figure 2), a collaboration and communication platform where contributors formed interest groups, shared discoveries, and exchanged ideas.
- **Linkurious** (Figure 3), a visualization system that provided visual graphs of the relationships between entities mentioned in the leaked documents.

From this survey and interview data, we identify several key design decisions and deployment strategies that appear to have contributed to the security successes of the project.

[1] All screenshots were approved for publication by ICIJ.

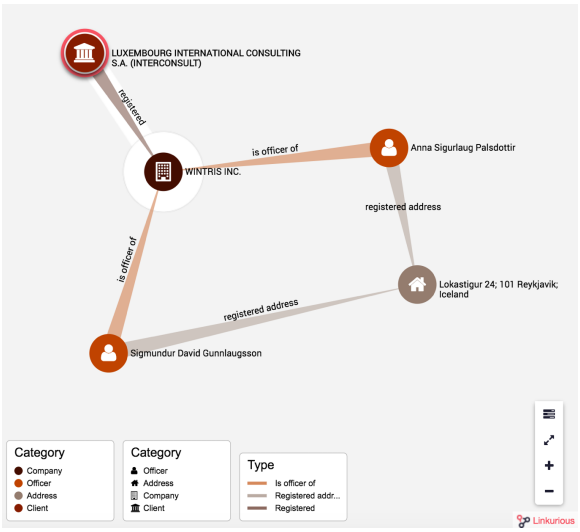

Figure 3: **Linkurious.** Screenshot of the system that visualizes links between entities mentioned in the Panama Papers documents. *Courtesy: ICIJ.*

For example, we were surprised to learn that project leaders were able to consistently enforce strict security requirements–such as two-factor authentication and the use of PGP–despite the fact that few of the participants had previously used these technologies. Our findings suggest that journalists found the collaboration systems provided so *useful* that they relied on them in spite of sometimes onerous security requirements. We observe that project leaders also frequently and consistently articulated the importance of security measures, explicitly cultivating a sense of collaboration, mutual trust and shared security responsibility among system users. Moreover, this organizational buy-in for security measures went beyond rhetoric: in one instance, the organization bought and set up phones as second factors for journalists who did not have them.

From these and other findings, we distill lessons and recommendations for integrating computer security measures into future socio-technical systems. For example, we recommend normalizing secure communication requirements to reduce the decision-making burden they may otherwise impose. In the Panama Papers project, for example, making PGP a default tool and ensuring everyone had a PGP key meant that participants did not need to expend additional energy evaluating secure communication options. We also identify opportunities for future research, such as comparing this to other security successes to determine which factors are necessary and/or sufficient to achieve similarly effective secure socio-technical systems. Instrumenting technical systems to achieve a more complete picture of activity and possible compromises would also contribute to this understanding.

In summary, we make the following contributions:

- We analyze *quantitative survey data* from 118 journalists involved in the Panama Papers project.
- We conduct *semi-structured, in-depth interviews* with key stakeholders–including editorial and technical staff– involved in designing and implementing the technical systems used in the collaboration.
- From these two datasets, we investigate the *socio-technical systems* that supported the realization of the security goals of the Panama Papers' project.
- We identify an actively maintained and explicitly articulated culture of security that leveraged peer-oriented trust and accountability. We also identify several technical security issues that may have been present, but do not appear to have led to compromise in this case.
- Based on this case study, we make *recommendations* for future socio-technical systems with strong computer security requirements and identify opportunities for future work.

Overall, the Panama Papers project—which required international collaboration among hundreds of journalists over an entire year—is a unique case study that provides insight into the design and deployment of secure communication and collaboration systems for diverse, distributed organizations. To the best of our knowledge, this is one of the first in-depth case studies of such a security success. Though this paper is neither a comprehensive description of the technical features of the Panama Papers' systems, nor a prescription for technical security best practices, we believe the insights presented here–taken in conjunction with existing technical security best practices–provide a valuable foundation for secure collaborative system design.

2 Background

In this section, we provide specific background on the Panama Papers project (unless otherwise noted, details here are sourced from [12], published by ICIJ). Additional related work is discussion in Section 7.

The International Consortium of Investigative Journalists (ICIJ) is a non-profit, selective-membership organization founded in 1997. Comprised of just under 200 investigative journalists in more than 65 countries, since 2012 ICIJ has obtained several caches of leaked documents that have led to collaborative investigations across news organizations around the world (e.g., [28–30]). Yet, in the words of one ICIJ staffer interviewed for this paper, the Panama Papers project [31] — which lasted from approximately May 2015 to April 2016 — was where the organization's work collaborative and analytical systems "all came together."

Consisting of over 11.5 million documents in dozens of formats occupying 2.6 TB of disk space, the Panama Papers dataset was by far the largest and most complex that ICIJ had handled (the "Offshore Leaks" project, by contrast, comprised only 260 GB [13]). While just one staffer was devoted to research during ICIJ's first major leak project in 2011, by 2016, data and research positions comprised half of ICIJ's 12-person staff.

To deal with the enormous scale and complexity of the data, as well as facilitate the large, globally distributed team required to investigate it, ICIJ's Data and Research Unit built and/or refined several systems whose development had begun during prior document-based projects. Favoring open-source technologies, they chose Tesseract [7] to OCR the documents, Apache Tika [2] for document processing, and Apache Solr [1] for indexing. The UI for this last platform also became its namesake, Project Blacklight [6] (see Figure 1).

ICIJ also developed a secure communication hub–called Global I-Hub–by customizing OxWall [5], an existing open-source messaging platform (Figure 2). Finally, ICIJ licensed the Linkurious software [4] to visually graph relationships among entities that appeared in the data (Figure 3).

3 Methods

To better understand the decisions that shaped the Panama Papers' suite of collaboration systems–as well as identify factors that may have contributed to the successful maintenance of the group's security goals–we conducted two studies: an analysis of survey data collected from Panama Papers project contributors by the ICIJ, and a semi-structured, in-depth interview with each member of the ICIJ staff who had significant influence over the security features and policies related to the Blacklight, I-Hub and Linkurious systems.

3.1 Participants

All survey participants are investigative journalists who actively participated in the Panama Papers project. All interview participants currently work full-time for the ICIJ and/or had a significant role in determining the security features and requirements of the collaboration systems used throughout the project by the journalists surveyed. In the results presented here, participants completed either a survey *or* an interview.

Survey. Survey participants were 118 journalists working in 58 different countries representing every continent except Antarctica. No other demographic data was collected. This sample represents approximately 33% (118 of 354) of all non-ICIJ staff who worked on the project.

Interview. ICIJ consists of only twelve full-time employees. For this study we interviewed all five of the ICIJ personnel with significant editorial or technical input on the systems used during the Panama Papers project. In-

Security Practice	Unaware	Never	Few	Occasionally	Frequently
Passphrase	9%	21%	13%	15%	52%
Two-factor	16%	29%	14%	13%	42%
PGP	14%	34%	10%	17%	25%

Table 1: **Familiarity with and Usage of Security Practices Prior to Project (N=118).** Scale items were "Never heard of it before" (Unaware); "Knew about it, hadn't used" (Never); "Had used a few times" (Few); "Used occasionally" (Occasionally) and; "Used frequently" (Frequently).

terview participants were two technical and two editorial management staff of ICIJ, as well as the journalist who received the original Panama Papers materials and worked closely with ICIJ on the system requirements. Of these five participants, two participants were women and three were men. To maximize the insight gained from these interviews, we designed the interview script using information from a careful review of public information available about the systems (e.g., [10, 36]), as well as insight from an IRB-approved background (pilot) interview with an individual member of the Panama Papers project who had intimate knowledge of the systems involved. The team then collected and iteratively refined the major themes for the interviews, customizing their content based on the individual's primary (self-identified) role in the project as either an editorial (E) or technical (IT) leader.

3.2 Materials

Materials consisted of a survey and two interview scripts, described here and reproduced in Appendices A and B.

Survey Instrument. The survey was created by ICIJ to investigate collaborating journalists' use of the Blacklight, I-Hub, and Linkurious systems used during the Panama Papers project, as well as their experiences with the security of these systems. In this paper, we focus on the 10 survey questions related to the use and security of the systems provided to journalists by ICIJ (see Appendix A). In addition to these security-related questions, the survey also captured information about the value to journalists of other services provided by ICIJ.

Interview Scripts. We created two distinct, but mostly overlapping interview scripts for the editorial and technical interview participants. Topics for both groups included questions about the participants' background, their experience with the overall system, system functionality, any training they offered as part of the project, any breaches or failures they were aware of, and the potential scalability of the system. Additionally, we asked editors about how they selected and recruited journalists for project participation. Please see Appendix B for the complete interview scripts.

3.3 Procedure

Survey. The survey was conducted between July 28th and August 15th, 2016 by the ICIJ. Participants completed the survey via a Google form and took around 10 minutes to complete. Participants could choose to answer the survey anonymously or provide their name if they wished. ICIJ provided us with the survey responses as a de-identified dataset. Participants were not provided an incentive to take the survey.

Interview. We interviewed participants between December 2016 and January 2017. Interviews typically lasted about one hour and were conducted via telephone/online video/voice conference (four), with one taking place in person. All participants spoke fluent English and were interviewed in English. Participants were not provided an incentive to participate in the interview.

3.4 Data Preparation and Analysis

Once all interviews were complete, we transcribed the audio recordings, producing 96 pages of text. Using an inductive process we completed an initial round of qualitative coding to identify key themes, as substantive categories emerged from the data via grounded theory analysis [19]. These themes were then evaluated and refined through group discussion among all researchers, with a goal of capturing the core variables constituting our participants' experiences.

3.5 Ethical Considerations

Our entire protocol was IRB approved. Furthermore, because of the sensitive nature of our interview topic, we took extra precautions to maintain the privacy and anonymity of research participants. We explicitly did not request information about or publish details about security protocols that could compromise source identities, sensitive information, or future work.

All interview participants agreed to be audio recorded during the interview and answered all of the questions in the interview script. We stored and transmitted audio recordings only in encrypted form and used de-identified transcripts for the majority of the data analysis.

4 Results

In this section, we present results from the survey and semi-structured interviews.

4.1 Survey Results

Apart from de-identification, the survey data analyzed below is a summary of the un-redacted responses (n=118) and comments (n=57) from the 118 journalist contributors who completed the ICIJ survey. Where relevant, we have included representative comments alongside the survey results. We identify quotes using only a letter (J for journalist) and participant number.

4.1.1 Prior Familiarity with Security Practices

The challenge of meeting security goals when working with non-expert users has been widely documented (e.g. [8]). To evaluate the significance of the Panama Papers project as a "security success story," we analyzed survey results to determine whether prior security expertise of the journalist contributors may have been a factor.

In fact, in response to a question about prior familiarity with digital security practices (see Table 1), almost half of participants indicated that they were "Unaware" of or had "Never" used PGP or two-factor authentication prior to this project (47% and 45% respectively).

Familiarity with passphrases (i.e., passwords created by concatenating multiple dictionary words, along the lines of [52]) was somewhat greater, with only about a third (31%) reporting that they were "Unaware" of or had "Never" used a passphrase. More than half (52%) of participants reported that they frequently used a passphrase prior to participation in this project, while 42% reported they frequently used two-factor authentication. Only one-quarter (25%) reported that they frequently used PGP prior to participation in the Panama Papers project.

Given journalists' limited familiarity with strong security practices prior to the Panama Papers project, we note that ICIJ's decision to mandate PGP for all collaborators is especially striking. We discuss the implications of this further in Section 5.

4.1.2 Perceived Difficulty of Security Compliance

Each of the three primary systems journalists used for this project —Blacklight, I-Hub, and Linkurious—had a distinct login that required two-factor authentication for every sign-on. Moreover, every journalist on the project was required to use PGP for password-reset and some system notifications. Despite relatively limited prior exposure to some of these security practices, however, participants reported that they perceived it to be relatively easy to comply with these requirements.

On a seven-point scale from 1 ("Super easy") to 7 ("Extremely Hard"), participants' overall mean rating was 3.13 (see Table 2), with the majority (63%) rating compliance with the security requirements on the "easy" side of the scale. As one participant put it:

> I am kind of technologically challenged, so the fact that I was able to navigate these security features means it was probably as simple as it could be while still being effective. (J11)

Super Easy	1	15%
	2	31%
	3	17%
	4	14%
	5	13%
	6	7%
Extremely Hard	7	3%

Table 2: **Perceived Difficulty of Security Compliance (N=118).** On a scale from 1 - 7, where 1 is "Super easy" and 7 is "Extremely hard", how challenging was it to comply with the digital security requirements?

Meanwhile, only 10% of participants (12/118) rated the difficulty of complying with security practices as extremely hard ("7": 3% or "6": 7%).

Participants' low difficulty ratings of complying with these security requirements is especially surprising given that they include use of PGP, which prior work indicates is notoriously difficult to use (e.g., [40, 60]). We discuss possible explanations for these results—including the participants' trust in the team leading the project—in subsequent sections.

4.1.3 Perceived Utility of ICIJ Technology Services

Research indicates that motivation can play a significant role in the adoption of security practices in organizations (e.g., [23, 54]), and is increased if users find a system useful—or even necessary—to achieving their primary work objectives [57].

When rating the utility/necessity of the technology services provided by ICIJ (summarized in Table 3), the vast majority of participants reported that the technology was essential (83% for data and tools and 78% for coordination). Though less than half of participants (43%) reported that the training was essential, almost all participants (95%) rated the training as at least "useful." None of the 5% of journalist-contributors who did not find the training useful commented on the training, though others did comment specifically on their interest in additional training. For example:

> I would like to receive more training at digital security tools. It was really useful. I learned for myself how to encrypt my computer and find out how vulnerable was my information, due to my lack of expertise using digital security tools. (J81)

Service	Unnecessary	Not Useful	Useful	Very Useful	Essential
Data	0%	0%	4%	13%	83%
Coordination	0%	0%	8%	14%	78%
Tools	0%	0%	4%	13%	83%
Training	2%	3%	23%	29%	43%

Table 3: **Necessity and Usefulness of Technology Services Provided by ICIJ (N=118).**

Service	Never	Daily	Weekly	Monthly
Blacklight	0%	64%	33%	3%
I-Hub	3%	41%	48%	8%
Linkcurious	19%	4%	45%	31%

Table 4: **Frequency of Use of ICIJ Technologies (N=118).** Frequency of use during the three months preceding publication; "monthly" includes responses "every now and then."

4.1.4 Frequency of Use of ICIJ Technologies

In order to assess how well contributors' reported usefulness of the ICIJ systems matched their actual behaviors, we analyzed survey data on how frequently journalists used Blacklight, I-Hub, and Linkurious (see Figures 1-3). These results are summarized in Table 4.

The majority of respondents (64%) indicated that they used Blacklight—the document-search platform where contributing journalists could access the leaked documents—at least daily during the three months prior to the project launch date in April 2016. One third (33%) used Blacklight at least weekly, and only 3% used it monthly.

The vast majority of respondents (89%) used I-Hub—the collaboration and communication platform with forum and chat features—daily or weekly. Only 8% used I-Hub only monthly, while just 3% reported never having used it.

By contrast, a significant portion (19%) of respondents indicated they had never used Linkurious—the system that provided visual graphs of the relationships between entities mentioned in the leaked documents. About a third (31%) said they used it monthly and nearly a half said they used it weekly. Only 4% used it daily.

4.1.5 Collaboration Outside Home Organization

A key objective for ICIJ in facilitating the Panama Papers project was to encourage inter-organizational collaboration among participating journalists, to maximize the quality and impact of the resulting publications. The degree of collaboration therefore offers insight into both the *utility* and *usability* of these systems. Given the global distribution of the journalist-contributors, collaborative data management strategies like using local-only servers or in-person meetings, were not feasible. These circumstances therefore also gave rise to specific *technical* security requirements for inter-organizational collaboration.

Survey participants were asked to rate how much they collaborated outside of their own organization on a scale of 1 - 7 (where 1 indicated "I worked independently" and 7 indicated "I've collaborated more than ever"). Nearly one third (32%) of participants indicated they had collaborated with journalists outside their organization "more than ever" during the Panama Papers project, and the vast majority of participants (74%) responded on the positive side of the scale (5, 6, or 7), with a mean rating of 5.33. Only 13% indicating lower levels of inter-organizational collaboration by responding on the negative side of the scale (1, 2, or 3). This data is summarized in Table 5.

4.1.6 Contributor Suggestions about Security

The survey data we analyzed also included one open-ended question: "Do you have any suggestions or comments about the digital security tools and requirements for this project?" Fifty-seven contributors offered open-ended feedback. While the themes of these comments varied, the most frequent theme was a feature request (14% total, of which more than half were requests for additional security features). The second most common themes were compliments (13%), statements affirming the need for security (5%), and requests for additional training (4%). Notably, only 3% of comments described the project's security requirements as a barrier to work.

For example, several participants (5) specifically mentioned issues around phone-based authentication.

> The Google Authentificator [sic] tool... when I changed my phone (twice during the investigation) I had to communicate with the support team to reboot the passwords. (J118)

> At certain times security turned into a barrier into getting more done... Every time a cellphone died or went missing (frequently) I needed to reconfigure authentication. (J68)

However, another participant noted that while security was a barrier, it was worth the slow-down:

> It's always a pain and even slowed us down. But this work is important and anything to keep it secure is fine. (J78)

Finally, others explicitly called-out the need for security and even praised ICIJ's focus on it:

> I like the fact that ICIJ considers security as a priority. Maybe ICIJ can explore other ways to

Independent	1	3%
	2	4%
	3	6%
	4	14%
	5	25%
	6	17%
Collaborative	7	32%

Table 5: **Collaboration Outside Own Organization (N=118).** Scale items were "I worked independently" (1) and "I've collaborated more than ever" (7).

> find log in ways that will not discourage potential users while at the same time putting security of our work a priority. (J109)
>
> Not been an expert, I believe the ICIJ team has done a fantastic work on security. (J111)

As we will discuss further in Section 5, this trust placed by contributors in the ICIJ team likely contributed to that team's ability to mandate security requirements.

4.2 Results from Interviews

We now turn to a discussion of our interview results, according to the topics from the interview script. Where relevant, we include verbatim quotes from participants to illustrate our findings. We identify quotes using only a letter (IT for technical staff, E for editorial staff, including the journalist who originally received the Panama Papers documents) and participant number.

4.2.1 Security Goals and Threat Model

Because the documents at the center of the Panama Papers project related largely to tax evasion, government actors—who could expect to recoup lost revenue through their exposure—were explicitly *not* considered to be part of the threat model for the project. That said, the companies, criminals (such as tax evaders, money-launderers and drug-traffickers), and politicians who were implicated in the documents were all identified as actors who could potentially confiscate locally-held data as well as threaten, imprison ,or even kill the journalists involved and/or block publication or access to the work. Given the size and resources of ICIJ, the primary security goals prior to publication therefore centered on protecting the source of the documents, maintaining the secrecy of the project, and maintaining the availability of the Blacklight, I-Hub, and Linkurious systems.

While our research participants only explicitly mentioned DDoS attacks and inadvertent project exposure as risks, training documents provided by participants indicate a range of security concerns, such as: spyware/malware, network monitoring, weak passwords/password reuse, physical interception of data (via locally-stored, unencrypted data or printouts) and legal attacks via third-parties. For example, a training document explicitly warned contributors against using third-party applications to translate, OCR or visualize the Panama Papers data, and encouraged storing local data from the project only in encrypted, hidden volumes.

These concerns informed the system design in myriad ways. First, both the sheer volume of the data–and the goal of protecting its source–led in part to the decision to use a centralized, remotely-accessible method of sharing the documents, rather than providing contributors with individual hard drives, as ICIJ had done in previous projects. As one ICIJ staffer put it:

> This is sensitive data that has been leaked to ICIJ for a reason, and that those sources are trusting us with being... guardians of that information and protectors. So it's not for us to give away to anybody, not even a trusted colleague. (E2)

Instead, the centralized system allowed ICIJ to grant all journalist contributors *access* to the documents, while still allowing ICIJ to monitor–and restrict–the volume of data that they could download from the system.

Second, the lack of a nation-state adversary–in conjunction with the specifics of Amazon Cloud's contractual agreements–made cloud-hosting an option. It was also a technical requirement, due to the volume of data involved and the need for substantial pre-processing.

4.3 System Design

Informing and interacting with ICIJ's security goals for the project were the organization's driving journalistic objectives: supporting high-quality, high-impact reporting and publications. Due to the enormous volume of data and documents involved (2.6 TB consisting of about 11.5 million documents), as well as their global nature, remote search and collaboration were essential—priorities that were clearly shared by both the editorial and technical staff:

> The needs are... communicate, and search documents, and to do it collaboratively. (E3)
>
> One of the more important impacts was that journalists discovered how convenient, powerful and good it is to collaborate... I think that the I-Hub contributed to this: to teach them how to interact, and it is a really good thing to share knowledge, share documents, share data, and build these networks. (IT2)

One reason the multi-national collaboration was essential was the variety of formats and languages within the source material, especially since participants were warned–through training, tip-sheets, and regular messaging from project leaders–against using third-party tools like Google Translate due to security concerns. ICIJ's

tools were therefore crucial to effective collaboration across timezone and language barriers:

> With cultural barriers, with language barriers and with time zones and all that... I think it was just the speed and the friendliness... it made 11 million documents look easy, look doable, and look—because it was easy and friendly to use, it became addictive to the reporters doing searches... and I like that. (E2)

Indeed, explicitly cultivating collaboration was a key design goal of I-Hub in particular, and it seemed to work:

> You cannot collaborate on email, or encrypted email, or Signal. You need a real space that feels comfortable and friendly and it's colorful, and [I-Hub] was. (E2)

> The forum was never as used and crowded as this time... It felt like everybody was sharing [and] working very collaboratively. (E1)

One reason that I-Hub may have been so easily adopted was its explicit similarity to familiar technologies (see also Figure 2):

> You can upload files, you can "like" a topic. You know, which is something that we're all so used to in the Facebook world. But that simple kind of "liking" thing also helped reporters bond together and encourage one another. And they were not going crazy with the likes, you know, most of the time people were not "liking" things, they were actually contributing useful information. But sometimes, you know, when somebody has made an important discovery...it just helped tremendously with providing a sense of team. (E2)

As we discuss further in Section 5, the fact that the ICIJ explicitly cultivated and supported such a collaborative culture—and that this collaboration was core to the success of the project itself—helped lay the groundwork for users' acceptance of strong security requirements.

4.3.1 Selecting Journalist Collaborators

In line with prior research on investigative journalists (e.g., [46]), our survey results indicate that the majority of the journalist contributors to the Panama Papers project were not security experts. Since any member of the collaboration is a potential "weakest link", we examine how these collaborators were chosen.

While ICIJ explicitly sought project contributors based in as wide a range of countries as possible, the core group of journalists (which numbered approximately 100 as of September 2015 and grew to nearly 400 by project launch in April 2016) were all existing ICIJ members. Interestingly, members who brought in non-member colleagues were considered responsible for disseminating and enforcing security protocols set by ICIJ:

> We would reach out to our member and trusted person... then the trusted journalist talks to a very small group of people in his own media organization... And then, if they get assigned to do the story, then we would train them, we would give them access to platforms... It's up to the trusted member and reporter to enforce all the rules and regulations with any person that that reporter bring on board. (E2)

Our interviews suggest that explicitly leveraging trust relationships within an established social network helped maintain the project's security requirements even as new members joined. While in practice this resembles a "web of trust" model, we note that unlike some traditional web-based implementations, each human "link" in this chain had a strong-tie connection to their closest link.

4.3.2 System Security

We now turn to a discussion of the security decisions made in the design and maintenance of ICIJ's systems, based on our interviews.

Technical Security. Key security aspects of all systems includes careful vetting of the source documents (including scanning them for known malware), deploying well-tested HTTPS, and requiring two-factor authentication for each of the three core systems.

The team experimented with multiple versions of two-factor authentication, including virtual machines (discarded as too complex) and browser extensions (discarded as insufficiently secure). Eventually, they settled on a smartphone-based app solution, which proved scalable despite initial concerns:

> You have to have a smartphone. And, we had a little discussion about, "Is this going to work?" Because Africa is big on cell phones, but mostly they're not smartphones... And then, when we started adding partners to the Panama Papers, everybody pretty much ended up having a smartphone. (E3)

Secure Defaults. One striking security decision was making PGP-encrypted email the default communication method for essential system functions. By summer 2016, participants were required provide a PGP key in order to obtain system credentials (including reset/recovery). Today, all notification emails from those systems are also encrypted by default.

Initially, however, contributors could receive password rest information via HushMail HushMail, and unencrypted system notifications still included details like

the summary of an updated thread. As security concerns increased post-launch, however, all details were eliminated from notifications until default PGP encryption could be implemented. Yet we note that the security culture among these journalists was strong enough by this point that they were willing to tolerate several months of *reduced functionality* for security purposes.

Achieving these secure defaults, however, was not the result of voluntary collective action: at some point, ICIJ mandated that all contributors create and use a PGP key:

> It was not a choice... If somebody did not get themselves a PGP, he did not get access to the forum and to the I-Hub. (E1)

A helpful side-effect of this requirement, however, was that it became possible for PGP-encrypted email to become a default for communication even *beyond* the I-Hub—and it was, even for seemingly non-sensitive material. As one core editorial affiliate put it:

> We had a rule in our team that whatever is about the Panama Papers—and if it's only about, I don't know, "Let's meet at nine, okay?" then we encrypt it because we encrypt everything that has to do with the Panama Papers. So that was our rule... the automatic step was to encrypt. (E1)

By creating secure defaults—especially ones that were useful outside of the project's infrastructure—the security achieved *within* the Panama Papers project systems also enhanced journalists' level of security *beyond* them.

Human Support and Communication with Users. Both technical and editorial staff emphasized the incremental was in which security features were rolled out. Moreover, they highlighted that security mandates from ICIJ were counterbalanced by increasing user investment in the systems, supported in part by open feedback channels and the addition of user-requested features:

> I said, "If you have any suggestions or any questions regarding the platforms, email me." (E3)

> We also encouraged the community to tell us through the Global I-Hub. There was a group called "data geeks" or something like that, and we encouraged them to tell us where we could improve. (E3)

ICIJ also provided accessible (human) technical support:

> We also have a support channel... So we're always assisting them all the time with their technological needs... Some of them forgot to change their phones... [and] didn't know how to re-install or how to reconnect with a new authenticator. (E2)

The result was a pace of security upgrades that matched users' investment in and need for the systems:

> So we have people to teach them how to [set up their PGP key], we have a support team that can help them... It went well because they were interested in keeping the access to the [platform]. (IT2)

In addition, these open lines of communication led to broad-based improvements in the platforms' functionality. For example, the user-suggested functionality of "batch search" was mentioned by four out five interview participants as one of the most valuable features of the Blacklight system:

> I was very glad that we could do batch searches in the end, which is a huge help. (E1)

Security Disagreements. Of course, security-related disagreements did arise. As security concerns increased post-launch, for example, reliance on the more usable Hushmail was scrapped in favor of PGP:

> It's much easier to create a Hushmail account. It's like creating a Google account. You know, like it takes that long [snaps finger]. Like nothing. I think that they say it's an encrypted system end-to-end and other things, but the reality is also that you don't know. (E2)

> At one point, I approached my managers and I said, look, everyone has serious doubts about HushMail... we just need to change our policy. (IT1)

ICIJ technologists also considered using CryptDB [3, 49], to encrypt the source documents while keeping them searchable. Yet while both primary technologists agreed that CryptDB was not a good fit, their reasoning around this decision was different. While one participant cited a mismatch in threat model, another had concerns about CryptDB's maturity:

> I don't think that there is any benefit in encrypting data at rest. We had this discussion early on in the project. One of the proposals was to use an encrypted version of MySQL [CryptDB]... the passwords have to be stored on the servers themselves... So what's the point? (IT1)

> We tried to use CryptDB, which is an encrypted database, but it was a new project and it didn't work... because the project was not stable enough. (IT2)

4.3.3 Security Weaknesses

Incidents. Our interviewees knew of no system breaches that took place during the course of the Panama Papers

project. Prior to launch, there was only one occasion when system monitoring suggested a possible attack:

> We had to ask one of our partners to bring his computer because we were detecting some weird requests to our systems. (IT2)

Once the partner in question changed machines, the requests stopped, though the underlying cause was never determined (the device was reviewed by the partner's organization, but no report was made to ICIJ).

Two security incidents occurred post-launch, both centering on the exposure of the systems' URLs, which had been intentionally kept secret. Due in part to the cost of more robust DDoS protection, project leads opted to maintain endpoint secrecy:

> If someone gets the location of the servers, they can do several attacks... We are prepared for this, for brute force [authentication] attacks... But yeah they also can send a DDoS attack, for example... So we have to protect the location, the server location. (IT2)

However, this "security by obscurity" approach suffered from an accidental leak:

> For example, we have requested that no URLs were ever shared or showed on television, like URLs of our platforms. And [partner organization] forgot about it and shared URLs on television... When this thing happened with the URL, we had to basically disconnect everyone from the platform and change the URL. (E2)

Though exposure of the URL only enabled attacks on system availability/uptime—knowledge of the URLs alone did not provide access to sensitive data—ICIJ was concerned enough about these exposures that they chose to take the systems temporarily offline in order to change their locations.

Technical Limitations. Though ICIJ and its collaborators were able to maintain the project's security goals, our study suggests several potential technical security limitations in their approaches.

For example, while ICIJ focused heavily on preventative security measures (e.g., ensuring encrypted communications), systematic approaches to dealing with potential security incidents seemed limited. While some networking monitoring and logging was available for network activity and document downloads, no systematic approaches to detecting or responding to potential data exfiltration events or other system breaches were described by our interview participants. For example, the discovery of an accidentally broadcast system URL was handled in an apparently ad-hoc way.

We also observe a strong focus on communications security (e.g., PGP) but less focus on endpoint security. While ICIJ was in a position to mandate security measures around communications, their influence on endpoint and operational security was limited to occasional training opportunities and "best practices" documents shared with contributors that addressed password management, third-party tool use, use of new/unfamiliar networks and basic threat modeling. They also recommended (and provided instructions for) creating encrypted hidden volumes for project documents stored locally. However, we do not know of any measures taken to verify adherence to these guidelines by participants.

ICIJ may have deployed additional security measures that we did not learn about in our interviews, but we highlight these potential weaknesses to provide context for the overall success of the project. We encourage future system designers to take the lessons from this paper in conjunction with existing security best practices.

4.4 Results Summary

In summary, we found that a large group of geographically and culturally diverse journalistswere able to collaborate securely over roughly a one-year period. To achieve their security goals, they relied on established security mechanisms such as PGP and two-factor authentication, as well as less systematized security practices like a social-network approach to adding members.

Overall, our survey results suggest that participants felt that complying with the security requirements of these systems was relatively easy, in spite of the fact that a large proportion of them had never used security technologies including two-factor authentication and PGP prior to the project. This is even more striking given that the vast majority of participants reported using the Blacklight and I-Hub systems daily in the 3 months prior to the project launch, each of which required a separate, two-factor login for every sign-on.

Our interviews, meanwhile, offer insight into both the core system requirements of the Panama Papers project, as well as the specific ways—such as strong HTTPS, two-factor authentication, a PGP/encrypted email default, and centralized control of the documents—the project's security goals were met. Through secure defaults and strong trust relationships reinforced through these collaborative systems, the limited security incidents were well-tolerated and compromised none of ICIJ's major security goals for the project.

5 Discussion

We now step back and reflect on the contributing factors to the Panama Papers project's security success, and reflect on how these factors may usefully inform the design of secure journalistic collaborations, as well as usably secure socio-technical systems more generally.

5.1 Factors for Success

Useful and necessary system functionality allowed for security mandates. A key factor in the success of ICIJ's approach was that journalists found their systems both *useful and necessary*, independent of their security properties. Journalists needed these systems for their core functionality (i.e., access to the source documents and collaboration with their peers), making strong security requirements (such as two-factor authentication and PGP) acceptable trade-offs to gain and maintain access. ICIJ staff were aware of this dynamic:

> You have to keep a balance between functionality and security. Because if you start with all the good practices in security at once, journalists will react to that and they will resist it, because it will be less functional. (IT2)

Our findings here align with research from management science, such as the Technology Acceptance Model [15, 56, 57], which argues that successful technology adoption in organizations depends not on mandated compliance, but rather on (1) usefulness and (2) ease of use. These factors a blend of both "social influence processes" (e.g., working norm, voluntariness, and image) and "cognitive instrumental processes" (e.g., job relevance, output quality, perceived ease of use) [57]. Among these, however, "usefulness" (defined as the user's perception that the new technology will enhance their job performance and output) was found to be the most powerful determinant of user acceptance.

Normalized security practices and secure defaults. The Panama Papers project leads actively cultivated a security-conscious culture in which secure communications were the norm. This norm helped project participants avoid the need to make granular decisions about which interactions warranted secure treatment. Several of our interview participants clearly identified the value of this approach. For example:

> In this project we just routinely encrypted everything we wrote... Because we were just used to doing it and that helped us a lot as a team, that we understood that it's not such a big thing, it's not such a pain in the ass—but you're always on the safe side of it. (E1)

By contrast, prior work [18] on email encryption adoption in an activist organization identified issues around encryption of non-sensitive messages. By universally encrypting *all* project-related communication, the Panama Papers team avoided such social complexities.

Usable alternatives for secure communication minimized workarounds. The ICIJ's systems supported multiple forms of secure communication, giving users flexibility depending on their needs and task. For example, I-Hub enabled secure group communication:

> For colleagues who are not that experienced with PGP or Signal or whatever...[the I-Hub is] a good way to write secure emails or messages to each other. (E1)

Where ICIJ systems didn't meet a particular need, however, contributors often reached for tools mandated by ICIJ (e.g., PGP) or other secure alternatives, thanks to the overarching security culture of the project, and the familiarity with and trust in these tools that the project provided:

> I don't like using PGP on the cell phone particularly. So then I would mostly switch to other channels, like Signal. (E1)

System designers, meanwhile, were conscious of users' primary task objectives and strove to minimize the friction of security security processes:

> It had to be as secure as possible, and still allow working with it without doing a three-day procedure to get into the I-Hub. (E1)

Cultivating mutual respect and reciprocity. The Panama Papers project systems were the product of an iterative design process within a particular community (journalists) and use case (i.e., facilitating global collaboration around a large trove of documents). This gave the ICIJ team confidence that the systems honored both their needs and values as an organization, and those of the journalist-contributors:

> It's great, it's just software that is designed for journalists... and that's all we care about. (E2)

> Panama Papers is *the* project where we tried to apply all the lessons learned from the previous projects. (E3)

ICIJ also maintained a careful balance between mandating security protocols and adding user-requested features (e.g. batch search), creating a sense of balance and equal partnership between the organization and journalist-collaborators:

> Once you have users, users will ask for things. They're helpful, you know? So, batch searching feature, I did not plan that. But people started asking "Would it be possible?" And it's like, "Ah, sure. This is a great idea." (E3)

This culture of mutual interest and respect helped users accept–and even support–ICIJ's strong security requirements.

Consultation with security experts. The ICIJ team chose third-party services carefully, based on advice from outside security experts:

> In the beginning we talked a lot to security experts. We did not really tell them what we had, of course not. But we needed to know more about the whole issue and the [organization] explained a lot about it worked... and why it's secure... So I know this seemed to make sense, and we spoke to other experts and they said "Yeah, you're on the safe side with that." (E1)

For example, while there were initial questions about using cloud hosting, Amazon Cloud Services' contract promises to inform customers of government access requests, allaying some fears:

> Amazon has quite a good reputation when it comes to ensuring the confidentiality of the customers...Their policy is to inform organizations if a state agency has requested a form or information from them. (IT2)

Hushmail, on the other hand, was eventually abandoned due to uncertainty around its security properties:

> I dont even know whether [Hushmail] has end-to-end encryption. It's just completely...non-transparent. It's much better to use PGP. (IT1)

Although the technical security measures deployed by the ICIJ were not necessarily complete, we note that they were thoughtfully constructed. We encourage future system designers to similarly engage security experts and/or rely on current security best practices as much as possible during the design process.

Leveraging social relationships to build trust and shared responsibility. Strong trust and social relationships were integral to the Panama Papers' collaboration from the start: the initial group of contributers were all ICIJ members, and becoming a member requires the explicit support of multiple existing members.

In addition to leveraging their strong ties with existing members, actively cultivating a collaborative, trust-based ecosystem among the non-ICIJ journalists helped security practices permeate the otherwise disparate and physically disconnected group. In addition, frequent project updates and security reminders from the ICIJ team–as well as specific design elements of I-Hub–helped further develop this sense of team and trust:

> [On the I-Hub,] the small things, like the fact that there's an avatar and you can see the face of the journalist, and you can have direct communications and all that... it helps with trust. It helps with bonding. (E2)

This observation echoes prior work [37] which found that users make security decisions motivated in part by a desire to maintain social work relationships. Indeed, social pressure can nudge users towards security compliance even if that compliance is burdensome or time intensive [48]. Prior work [33] has also found that a high rate of in-group communication fosters greater trust. Our findings suggest that these factors all played in a role in the security success of the Panama Papers project.

Sustained emphasis on security. Project leads at ICIJ also clearly and frequently communicated the importance of security and what was at stake:

> In every editorial note I would write, I would remind [contributors] about some security measure, how it takes one of us to make a mistake for the whole thing to basically fall to hell, and you would lose an entire year of work, and we would be—a joke basically. Nobody would ever come to us again with any confidential information. So, I would remind them so they didn't feel comfortable and too confident. (E2)

Organizational resource commitment. A key success factor was the ICIJ's willingness to commit resources to developing useful and secure systems:

> [Collaborating] requires a team, and it requires systematic work... If there's no compelling need, journalists are not going to use it... It has to be enforced also by the managers and embraced by everyone. (E2)

Though stakeholders sometimes disagreed, developers actively sought management buy-in for creating long-term security solutions. For example:

> There is a tendency... to have this kind of quick solution and where it puts the load of the problem onto staff. The solution my managers proposed [for password reset issues].. created a huge support burden... Selling [long-term technical solutions] is a little difficult to directors... But when you do implement it, it works beautifully I think, and becomes an example to other organizations. (IT2)

5.2 Lessons and Opportunities for Future Research

For the computer security research community, this case study represent a rare example of security success, achieved despite many complicating factors. Examples include: mandating important but notoriously inconvenient and/or hard to use protocols, like PGP [60]; contributors' lack of prior experience with the mandated security practices; participants' wide geographic distribution and diverse native languages. Yet ICIJ was able to mandate their security requirements, and hundreds of contributing journalists adhered to—and even applauded—those requirements, allowing the project's security goals to be met. While the systems used in the Panama Papers project are not appropriate for every project, organization, or security scenario, we believe

this example offers important insights for those wishing to design similarly effective systems, especially from a human-centered perspective.

Lessons for Journalistic Collaborations. A key factor in the Panama Papers' security success was the reputation ICIJ had built for exclusive, high-impact investigations. Journalists approached for the Panama Papers' project were thus strongly incentivized to meet ICIJ's security requirements, which were required to gain access to the systems and the documents they held. The risk of being left out of future projects or ostracized by colleagues and partner organizations made the cost of security non-compliance particularly high. Similarly, the importance of clear communication around security suggests that tying security requirements to demonstrable professional advantage, along with clear expectation-setting (including negative consequences) are key factors in motivating journalists to adopt and maintain even potentially onerous security practices.

Recommendations for Socio-Technical Systems. Prior work shows that employees will often sidestep security requirements to focus on their primary tasks [24]. As the Panama Papers project demonstrates, however, when security measures are *integral* to those tasks, they may be better honored by users. This suggests that security measures perceived as a "bolt-on" to existing systems—especially if organizational leaders are not vocal about their importance—may engender avoidance behaviors from users. Similarly, insufficient attention by system and security experts to the specific work needs and task priorities of users may lead to brittle systems: tools and protocols that do not offer multiple methods for meeting a particular security requirement (e.g. text-based communication), may lead users to rely on insecure workarounds to meet their needs. This affirms prior work (e.g., [48, 55]) suggesting that ongoing attention to both security and primary work objectives by organizational leaders *and* security experts is key to creating and maintaining secure collaborative systems.

Opportunities for Future Work. Though our work has identified multiple factors that may have contributed to the effective security of the Panama Papers collaboration, we do not know which of these factors were *necessary*, nor which combination of them would have been *sufficient*. We also cannot tease out the importance of other potentially relevant factors, such as whether the small size of ICIJ itself helped facilitate organizational consensus on security issues.

Two key directions for future work, then, include (1) conducting additional case studies of socio-technical security successes and (2) comparing these case studies to clarify which factors are necessary and/or sufficient. While our findings support prior work on the value of social relationships for motivating security behaviors, exploring other motivations (such as professional norms or organizational identity) may highlight additional paths towards similar types of security success.

6 Limitations

The Panama Papers project provides a remarkable example of a diverse, highly-distributed group of journalists meeting the security goals of the coordinating organization. However, we know that no system is perfectly secure, and that even systems that appear to meet their security goals may have been breached. In this case, a highly-motivated and/or -resourced attacker could—without the organization's awareness—have potentially or actually compromised the systems we described here. We do not claim causality, ultimate system security, or lack of vulnerabilities, but rather identify factors that may have contributed to the ICIJ's success in achieving their security goals (protecting the source and preserving the secrecy of the project until the desired launch date) in a complex socio-technical system.

Thus, the measures describes above should not be interpreted as a guarantee of security or recipe for success, nor a complete technical description of the systems used. Indeed, we highlighted several technical limitations of the system and encourage readers to treat this case study as a potential starting point from which to incorporate other technical security best practices (e.g., mechanisms for detecting compromise or strengthening endpoint security). We leave a technical analysis of these still-evolving systems to future research.

Finally, because the survey instrument was designed by ICIJ, we could not control what questions were asked and how. We include the survey instrument in Appendix A for transparency.

7 Related Work

To the best of our knowledge, this paper represents one of the first in-depth studies of a security success story. Due to the novelty of such a case study in the security literature, below we examine related work in adjacent fields.

7.1 Security for Journalists and Activists

Recent work has studied computer security for journalists specifically, both individually [40, 44] and organizationally [45]. These works identified computer security challenges due to, e.g., the fragility of journalists' relationships with their sources, as well as the limited resources available within journalistic organizations.

Like NGOs and activist groups, journalists' work makes them high-value targets for cyberattack and surveillance (e.g., [20]). Certain nation-states have been known to monitor these groups and scan for evidence

of political dissent, by "eavesdropping, stealing information, and/or unmasking anonymous users" [43]. In addition to surveillance, such groups have also been the target of malware attacks and tailored phishing attacks, on which several case studies have been published [21, 39].

7.2 Security in Organizations

More generally, when considering computer security within organizations or other networks, users are often considered the "weakest link" [8]—a theme that has become common in a range of fields (e.g., [26, 48, 50]).

Usability studies have begun to amend this assumption, looking at how to strike the balance between security and usability (e.g., [35]). Work in this field shows that users make decisions informed by a rational concern for efficiency, so much so that many deliberately ignore security advice and training [24, 25].

Scholars have found that organizational culture is a critical component for the successful implementation of security policy [58]. For example, Kirlappos and Sasse [37] show that social relationships between employees impact compliance with security mandates. Blythe et al. [10] identified factors contributing to employees' security behaviors, including security knowledge and perceptions of responsibility. Thomson et al. [55] highlight the importance of integrating security awareness into an organization's daily culture. Pfleeger et al. [48] discuss the rollout of security mandates in the context of employees' mental workload and interaction with their primary task flow, All of these factors from prior work—peer trust relationships, organizational security culture and norms, and integration with primary tasks—are echoed in our findings.

Other fields, including managerial and behavioral studies as well as social psychology and sensemaking, also consider the role of employee-culture in general managerial compliance. Organizational culture, in particular, has been found to exert outsized influence on employee behavior [16, 22, 26, 34, 41, 48, 51, 53].

7.3 Security on Distributed Teams

As technology has enabled geographically distributed teamwork, top-down management has given way to decentralization and flat hierarchies [14] This change has security implications: top-down enforcement has been shown to be less effective than socially embedded, trust-based cultural compliance [37]. Moreover, top-down mandates can actually lead to employees' distrust of the organization [59] or harm productivity [27]. Our findings here—where security mandates were accepted and even supported by journalist-contributors—suggest that this distrust effect may be overcome by sufficiently strong social relationships and/or respect for the organization.

For digital rather than physical collaborations, computer security becomes critically important, and knowledge management in such teams is a topic of interest for researchers [9, 11, 38, 47]. However, with some notable exceptions [14, 33], the specific requirements of such teams for security compliance are understudied. Our research helps address this gap in the literature.

8 Conclusion

In this paper, we have explored a security success story: the case of the year-long Panama Papers project collaboration among hundreds of journalists around the world. We presented and analyzed survey data from 118 journalists involved with the project, as well as interviews with the editorial and technical staff behind the design and implementation of the collaboration tools used during the project. From these datasets, we distilled success factors and recommendations for designing and implementing secure socio-technical systems.

We found that users will accept strict security requirements in order to use tools critical to their core (non-security) efforts; that a strategy of reducing security decisions by making secure behavior the default and providing secure alternatives for functionality not directly supported may discourage insecure "workaround" behaviors; that leveraging peer relationships can help foster a collaborative culture with a shared sense of security responsibility; and that inviting—and engaging—input from users helps establish a sense of reciprocity that facilitates their adoption of security mandates. This case study demonstrates not only *that* meeting significant security goals is possible in a complex socio-technical system, but provides valuable insights into how similarly successful future systems can be designed.

Acknowledgements

We are especially grateful to our interview participants and the ICIJ Data Team for providing us access to the survey data and images of their systems. We thank undergraduate research assistants Brian Justice and Duyen Nguyen at Clemson for help transcribing the interviews. Finally, we thank our anonymous reviewers and our shepherd, Adrienne Porter Felt, for valuable feedback on an earlier version. This work is supported in part by the National Science Foundation under Awards CNS-1513575, CNS-1513875, and CNS-1513663.

Author Contributions

SM recruited participants, collected all data, and coordinated the writing and editing process. KC originated the study idea and EW and KC conducted data analysis. MA and EW contributed the literature review. MA helped prepare study materials and processed raw data. FR, SM and KC drafted sections of the paper and FR guided its framing for the USENIX Security audience. All au-

thors identified relevant themes and illustrative quotes, contributed to the discussion section, and reviewed and edited the final manuscript.

References

[1] Apache Solr. http://lucene.apache.org/solr/.

[2] Apache Tika. https://tika.apache.org/.

[3] CryptDB. http://css.csail.mit.edu/cryptdb/.

[4] Linkurious. http://linkurio.us/.

[5] Oxwall. https://www.oxwall.com/.

[6] Project Blacklight. http://projectblacklight.org/.

[7] Tesseract. https://github.com/tesseract-ocr.

[8] ADAMS, A., AND SASSE, M. A. Users are not the enemy. *Communications of the ACM 42*, 12 (1999), 40–46.

[9] ALAVI, M., AND TIWANA, A. Knowledge integration in virtual teams: The potential role of KMS. *Journal of the American Society for Information Science and Technology 53*, 12 (2002), 1029–1037.

[10] BLYTHE, J. M., COVENTRY, L., AND LITTLE, L. Unpacking security policy compliance: The motivators and barriers of employees' security behaviors. In *11th Symposium On Usable Privacy and Security (SOUPS)* (2015), pp. 103–122.

[11] BODEN, A., AVRAM, G., BANNON, L., AND WULF, V. Knowledge management in distributed software development teams: Does culture matter? In *4th IEEE International Conference on Global Software Engineering* (2009), IEEE, pp. 18–27.

[12] CABRA, M., AND KISSANE, E. Wrangling 2.6TB of data: The people and the technology behind the Panama Papers, 2016. https://panamapapers.icij.org/blog/20160425-data-tech-team-ICIJ.html.

[13] CAMPBELL, D. Offshore secrets: unravelling a complex package of data. *The Guardian* (2013). https://www.theguardian.com/uk/2013/apr/04/offshore-secrets-data-emails-icij.

[14] DAMM, D., AND SCHINDLER, M. Security issues of a knowledge medium for distributed project work. *International Journal of Project Management 20*, 1 (2002), 37–47.

[15] DAVIS, F. D., BAGOZZI, R. P., AND WARSHAW, P. R. User acceptance of computer technology: A comparison of two theoretical models. *Management science 35*, 8 (1989), 982–1003.

[16] DOUGLAS, P. C., DAVIDSON, R. A., AND SCHWARTZ, B. N. The effect of organizational culture and ethical orientation on accountants' ethical judgments. *Journal of Business Ethics 34*, 2 (2001), 101–121.

[17] ERLANGER, S., CASTLE, S., AND GLADSTONE, R. Iceland's prime minister steps down amid Panama Papers scandal, April 6, 2016. https://www.nytimes.com/2016/04/06/world/europe/panama-papers-iceland.html.

[18] GAW, S., FELTEN, E. W., AND FERNANDEZ-KELLY, P. Secrecy, flagging, and paranoia: Adoption criteria in encrypted email. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2006), ACM, pp. 591–600.

[19] GLASER, B. G., AND STRAUSS, A. L. *The Discovery of Grounded Theory: Strategies for Qualitative Research*. Aldine Publishing Company, Chicago, 1967.

[20] GREENWALD, G. *No Place To Hide: Edward Snowden, the NSA, and the U.S. Surveillance State*. Metropolitan Books, 2014.

[21] HARDY, S., CRETE-NISHIHATA, M., KLEEMOLA, K., SENFT, A., SONNE, B., WISEMAN, G., GILL, P., AND DEIBERT, R. J. Targeted threat index: Characterizing and quantifying politically-motivated targeted malware. In *23rd USENIX Security Symposium* (2014).

[22] HARRIS, S. G. Organizational culture and individual sensemaking: A schema-based perspective. *Organization Science 5*, 3 (1994), 309–321.

[23] HERATH, T., AND RAO, H. R. Protection motivation and deterrence: a framework for security policy compliance in organisations. *European Journal of Information Systems 18*, 2 (2009), 106–125.

[24] HERLEY, C. So long, and no thanks for the externalities: The rational rejection of security advice by users. In *Proceedings of the New Security Paradigms Workshop* (2009), ACM.

[25] HERLEY, C. More is not the answer. *IEEE Security & Privacy 12*, 1 (2014), 14–19.

[26] HU, Q., DINEV, T., HART, P., AND COOKE, D. Managing employee compliance with information security policies: The critical role of top management and organizational culture. *Decision Sciences 43*, 4 (2012), 615–660.

[27] INGLESANT, P. G., AND SASSE, M. A. The true cost of unusable password policies: Password use in the wild. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2010), CHI '10.

[28] INTERNATIONAL CONSORTIUM OF INVESTIGATIVE JOURNALISTS. Secrecy for sale: Inside the global offshore money maze, 2013. https://www.icij.org/offshore.

[29] INTERNATIONAL CONSORTIUM OF INVESTIGATIVE JOURNALISTS. Luxembourg leaks: Global companies' secrets exposed, 2014. https://www.icij.org/project/luxembourg-leaks.

[30] INTERNATIONAL CONSORTIUM OF INVESTIGATIVE JOURNALISTS. Swiss leaks: Murky cash sheltered by bank secrecy, 2015. https://www.icij.org/project/swiss-leaks.

[31] INTERNATIONAL CONSORTIUM OF INVESTIGATIVE JOURNALISTS. The Panama Papers: Politicians, criminals, and the rogue industry that hides their cash, 2016. https://panamapapers.icij.org/.

[32] INTERNATIONAL CONSORTIUM OF INVESTIGATIVE JOURNALISTS. The Panama Papers - Reporting Partners, 2016. https://panamapapers.icij.org/pages/reporting_partners/.

[33] JARVENPAA, S. L., AND LEIDNER, D. E. Communication and trust in global virtual teams. *Journal of Computer-Mediated Communication 3*, 4 (1998).

[34] JONES, R. A., JIMMIESON, N. L., AND GRIFFITHS, A. The impact of organizational culture and reshaping capabilities on change implementation success: The mediating role of readiness for change. *Journal of Management Studies 42*, 2 (2005), 361–386.

[35] KAINDA, R., FLECHAIS, I., AND ROSCOE, A. Security and usability: Analysis and evaluation. In *International Conference on Availability, Reliability, and Security (ARES)* (2010), IEEE, pp. 275–282.

[36] KING, G. Best security practices: An overview. In *Proceedings of the 23rd National Information Systems Security Conference, Baltimore, Maryland, NIST* (2000).

[37] KIRLAPPOS, I., AND SASSE, M. A. What usable security really means: Trusting and engaging users. In *HCI International* (2014).

[38] KOTLARSKY, J., AND OSHRI, I. Social ties, knowledge sharing and successful collaboration in globally distributed system development projects. *European Journal of Information Systems 14*, 1 (2005), 37–48.

[39] LE BLOND, S., URITESC, A., GILBERT, C., CHUA, Z. L., SAXENA, P., AND KIRDA, E. A look at targeted attacks through the lense of an NGO. In *23rd USENIX Security Symposium* (2014).

[40] LERNER, A., ZENG, E., AND ROESNER, F. Confidante: Usable encrypted email – A case study with lawyers and journalists. In *IEEE European Symposium on Security and Privacy* (2017).

[41] LUND, D. B. Organizational culture and job satisfaction. *Journal of business & industrial marketing 18*, 3 (2003), 219–236.

[42] MALTBY, J., AND DAMON-FENG, G. The Panama Papers: The story so far, and what comes next, December 16, 2016. `https://www.law360.com/articles/874074/the-panama-papers-the-story-so-far-and-what-comes-next`.

[43] MARCZAK, W. R., SCOTT-RAILTON, J., MARQUIS-BOIRE, M., AND PAXSON, V. When governments hack opponents: A look at actors and technology. In *23rd USENIX Security Symposium* (2014).

[44] MCGREGOR, S. E., CHARTERS, P., HOLLIDAY, T., AND ROESNER, F. Investigating the computer security practices and needs of journalists. In *24th USENIX Security Symposium* (2015).

[45] MCGREGOR, S. E., ROESNER, F., AND CAINE, K. Individual versus organizational computer security and privacy concerns in journalism. *Proceedings on Privacy Enhancing Technologies 4* (2016), 1–18.

[46] MITCHELL, A., HOLCOMB, J., AND PURCELL, K. Investigative journalists and digital security: Perceptions of vulnerability and changes in behavior. Pew Research Center, Feb. 2015. `http://www.journalism.org/files/2015/02/PJ_InvestigativeJournalists_0205152.pdf`.

[47] OSHRI, I., VAN FENEMA, P., AND KOTLARSKY, J. Knowledge transfer in globally distributed teams: the role of transactive memory. *Information Systems Journal 18*, 6 (2008), 593–616.

[48] PFLEEGER, S. L., SASSE, M., AND FURNHAM, A. From weakest link to security hero: Transforming staff security behavior. *Journal of Homeland Security and Emergency Management 11*, 4 (2014).

[49] POPA, R. A., REDFIELD, C. M. S., ZELDOVICH, N., AND BALAKRISHNAN, H. Cryptdb: Protecting confidentiality with encrypted query processing. In *Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles* (New York, NY, USA, 2011), SOSP '11, ACM, pp. 85–100.

[50] SCHNEIER, B. *Secrets & Lies: Digital Security in a Networked World*. John Wiley & Sons Inc., 2000.

[51] SCHRODT, P. The relationship between organizational identification and organizational culture: Employee perceptions of culture and identification in a retail sales organization. *Communication Studies 53*, 2 (2002), 189–202.

[52] SHAY, R., KOMANDURI, S., DURITY, A. L., HUH, P. S., MAZUREK, M. L., SEGRETI, S. M., UR, B., BAUER, L., CHRISTIN, N., AND CRANOR, L. F. Can long passwords be secure and usable? In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2014), CHI '14, ACM, pp. 2927–2936.

[53] SMIRCICH, L. Concepts of culture and organizational analysis. *Administrative science quarterly* (1983), 339–358.

[54] STANTON, J., MASTRANGELO, P., STAM, K., AND JOLTON, J. Behavioral information security: Two end user survey studies of motivation and security practices. *AMCIS 2004 Proceedings* (2004), 175.

[55] THOMSON, K.-L., VON SOLMS, R., AND LOUW, L. Cultivating an organizational information security culture. *Computer Fraud & Security 2006*, 10 (2006), 7–11.

[56] VENKATESH, V. Determinants of perceived ease of use: Integrating control, intrinsic motivation, and emotion into the technology acceptance model. *Information Systems Research 11*, 4 (2000), 342–365.

[57] VENKATESH, V., AND DAVIS, F. D. A theoretical extension of the technology acceptance model: Four longitudinal field studies. *Management science 46*, 2 (2000), 186–204.

[58] VON SOLMS, B., AND VON SOLMS, R. The 10 deadly sins of information security management. *Computers & Security 23*, 5 (2004), 371–376.

[59] WEIRICH, D. *Persuasive password security*. PhD thesis, University College London, 2005.

[60] WHITTEN, A., AND TYGAR, J. D. Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0. In *Proceedings of the 8th USENIX Security Symposium* (1999).

A Appendix: Survey Instrument

This appendix contains the questions from ICIJ's survey of contributing journalists for which we received data.

A.1 ICIJ Journalist Survey

We want to know your opinion about the project platforms and your experience working on the project. It should take you 10 minutes. Your honest feedback will be important to make adjustments to future investigations and we will use your answers only for ICIJ internal purposes. You can answer the survey anonymously, although we appreciate if you tell us who you are. Thanks for helping us to improve global collaboration in journalism!

1. Name [short answer]
2. Country [short answer]
3. Media Outlet [short answer]
4. Email [short answer]

5. How much did you collaborate with others outside your organization for this project?
(I worked independently) 1 2 3 4 5 6 7 (I've collaborated more than ever)

6. How would you rate the services provided by ICIJ throughout this project?
For 6.A-C, the scale was: Unnecessary, Not useful, Useful, Very useful, Essential.
A. Project coordination
B. Digital tools (I-Hub, Blacklight, etc.)
C. Training (tools, data and digital security)

7. How did you find the coordination of the project?
(Poor) 1 2 3 4 5 6 7 (Excellent)

8. How often did you use _____ during the last three months before publication?
For 8.A-C, the scale was: Every day, Two or three times a week, Once a week, Once a month, Every now and then, I never used the service, Other: [short answer].
A. Blacklight
B. I-Hub
C. Linkcurious

9. Which digital security practices were you familiar with prior to working on this project?

For 9.A-C, the scale was: Never heard of it before, Knew about it but hadn't used, Had used a few times, Used occasionally, Used frequently.

A. Passphrases (instead of passwords)
B. Two-factor authentication (Google authenticator)
C. PGP encryption (for email)

10. Which improvements (if any) would you like to see in Blacklight? [short answer]

B Appendix: Interview Instruments

This appendix contains our interview script for ICIJ editorial personnel and for ICIJ technical staff. We note inline where the interview script differed between editorial and technical staff.

Background

1. What was your [editorial/technical] background and/or main area of responsibility for ICIJ prior to the start of the Panama Papers project?
2. Prior to the Panama Papers, had you worked on any other collaborative investigative projects at ICIJ, or any other organization? If so, can you tell us a little bit about how the Panama Papers differed from these earlier efforts?

Overall System Design

1. Were you directly involved in the [technical] design [and/or deployment] of the collaborative systems used during the Panama Papers to store and/or share the source documents? If so:
   (a) What did you feel were the most important features of the system in terms of functionality? What were the most significant challenges to including these features?
   (b) What did you feel were the most important features of the system in terms of security? What were the most significant challenges to including these features?
   (c) We understand that PGP was required to distribute at least some system credentials. Can you tell me a little bit about why PGP was selected, and how that requirement was communicated to users?
2. Were any of the technologists who worked on the projects not ICIJ employees? If so, how were they selected for involvement? Was their access to the design and/or implementation details of the project limited in any way?
3. To the extent that you are aware, how did the systems evolve over the course of its use during the Panama Papers project? Have they continued to change since the launch? In what ways?
4. From your perspective, what were the most successful aspects of the system design and deployment? What were the least successful? What surprised you the most about how the system was used?
5. *For technical staff only:* Were regular backups performed on the system? If so, how were backups initiated and carried out?
6. *For technical staff only:* Was content stored on the system generally encrypted at rest? If so, was there a mechanism for searching this content?

Recruitment and Participation

For editorial staff only:

1. How did journalists generally get involved in the Panama Papers project? Were they recruited, or did they reach out to ICIJ?
2. What was the general process for vetting individuals or organizations for participation? Was anyone ever rejected? Why?
3. Was there a group of people who were responsible for verifying the authenticity of received documents and information? If so, what type of process did they use?
4. As more information was received, how was it integrated into the system? Who was responsible for this, and how was the process determined?

General System Functionality: BlackLight and I-Hub

1. We understand that there were two primary systems used to manage the Panama Papers project: BlackLight and I-Hub. In your own words, you could describe each of these systems, both in terms of their functionality and how they were implemented?
2. Did journalists have separate logins to the two systems? To the best of your knowledge:
   (a) Were there specific password requirements (e.g., length, various characters, etc.)?
   (b) Was two-factor authentication required?
   (c) How could users change/reset passwords? Were regular password changes required?
3. *For editorial staff only:* Were users allowed to upload files to either system? If so, were there any system features included to scan or clean these files?
4. *For technical staff only:* Were users authorized to upload files to either system? If so, was there any service/feature embedded with the file server, to detect and clean malware when a file is uploaded?
5. *For editorial staff only:* If users had a difficulty with one of the systems, what resources were available to them? Was providing user support a significant consideration in the design of the system?

6. *For technical staff only:* If users had a difficulty with one of the systems, could they contact the IT team directly? If so, what was the mechanism? If not, what types of resources or protocol was made available for these users?

I-Hub

1. Did all journalist users have the same level of permissions on the system?
2. What type of user could create new "chat rooms" or threads? Could administrators see all of these, and/or remove content, if needed?
3. *For editorial staff only:* Were there any features that you would have liked to see included in the system, but that could not be integrated for technical reasons? What were they?
4. *For technical staff only:* What type of encryption was implemented on this system? Was it end-to-end (in the style of PGP or OTR) or client-to-server (e.g. HTTPS connection to platform)?

BlackLight

1. *For editorial staff only:* How did the BlackLight system work? Why was BlackLight selected as the base project from which to create the Panama Papers system? What features do you wish it had that it didn't?
2. *For technical staff only:* Why was BlackLight selected as the base project from which to create the Panama Papers system? Was it difficult to adapt or secure for use on this project? In what ways?

Listserv

1. How did communications on the listserv differ from those on I-Hub?
2. *For technical staff only:* What were the functional/security differences between I-Hub and the listserv?
3. Are you aware of any instances where the listserv was used inappropriately? If so, how was this addressed, and by whom?

Information Security Training

1. Who generally provided security training for journalists? Who designed the content of the trainings?
2. Did you provide or design any of these trainings? If so, please tell me a little bit about how they were delivered and what content they contained:
   (a) Were they "live" (e.g. streamed) or recorded? Why or why not?
   (b) Did they involve hands-on exercises? Why or why not?
   (c) Was there any type of evaluation/grading of participants? Could a "failing" grade limit access or require the training be taken again? Why or why not?
   (d) How many different trainings/topics did each user have to engage before being granted access to the systems?
3. What was the goal of providing these trainings? Do you feel they were successful? What would you change or do differently around training for a similar project in the future?

Security Breaches and System Failures

1. To what extent was keeping the online location (i.e. URL) of the project an important security concern?
2. Was there a specific protocol for taking the system offline due to errors, updates or security incidents? How were these communicated to the users of the system (if at all)?
3. *For editorial staff only:* Were there specific plans in place for detecting and/or handling system exposures or security incidents? How were the users and/or publications involved monitored, if at all? By whom?
4. *For technical staff only:* Were there specific plans in place for detecting and/or handling security incidents? For example, were there automated intrusion detection systems, or checks on the locations of system access?
5. Without revealing specifics that could compromise continued use of the system, can you share a general sense of what kind of security incidents happened during the project, and how they were handled?

Scaling and Future Development

1. Do you feel that you would use – or encourage others to use – this type of system for collaborative investigative projects in the future? Why or why not?
2. From both a functionality and support perspective, do you think the systems used for the Panama Papers are scalable to a larger number of projects and/or users?
3. Are there any [design or deployment / technical or system design] lessons you learned from this project that you intend to apply to the design of future systems, whether for similar projects or not? If so, what features or aspects would you keep or change for other projects, and why?
4. Would you change the content or mechanism of training or support for future systems?
5. Is there anything else about this project that you'd like to tell us or think we should know?