

User Study of Tor Relay Operators' Attitudes Towards Automatic Updates

Meitong Wang



Master of Science
Cyber Security, Privacy and Trust
School of Informatics
University of Edinburgh
2023

Abstract

User attitudes toward updates are increasingly important for usable security. Despite considerable technical advancements, understanding these attitudes remains crucial. This dissertation focuses on identifying the factors that shape user update decisions and comprehending the favorability of automatic updates, specifically in the context of Tor relays. This study holds significance as it aids developers in crafting automatic update mechanisms that account for the specific considerations of Tor relay operators. To achieve its research aim, this dissertation undertook a comprehensive literature review followed by an empirical study. The latter involved administering questionnaires to Tor relay operators. The research yielded several key findings. Existing literature confirmed the diverse factors influencing user update decisions. In addition, the empirical study unveiled specific factors unique to Tor relay operators, notably the need for control. Furthermore, the study highlighted concern among relay operators, particularly with security aspects, in relation to automatic updates. As such, this dissertation advocates for the design of automatic update mechanisms that prioritize security, entail low perceived costs, balance control, ensure effective communication, and maintain platform consistency.

Research Ethics Approval

This project obtained approval from the Informatics Research Ethics committee.

Ethics application number: 775342

Date when approval was obtained: 2023-05-27

In this project, data collection was conducted entirely online. The combined participants' information sheet and consent form is included in the appendix A.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified.

(Meitong Wang)

Acknowledgements

I would like to express my heartfelt gratitude to my supervisor for his invaluable guidance, unwavering support, and insightful feedback throughout the course of this research. His expertise and encouragement greatly contributed to the completion of this dissertation.

I extend my sincere thanks to the participants who took part in the survey, their willingness to share their insights and experiences played a pivotal role in shaping the outcomes of this study. I am also deeply appreciative of the Tor project leader who generously provided feedback during the design phase, enriching the quality of the research.

My appreciation goes to my parents and friends for their continuous encouragement and belief in my abilities. Their emotional support has been a source of strength, and I am truly grateful for their presence in my journey.

Last but not least, I want to acknowledge my own determination and resilience. There were challenging moments along the way, but I am thankful for my perseverance which kept me moving forward towards the successful completion of this dissertation.

Table of Contents

1	Introduction	1
1.1	Motivation	1
1.1.1	Tor and Relay Operators	1
1.1.2	Automatic Relay Update Framework	2
1.2	Objectives	2
1.3	Results and Outcomes	3
1.4	Structure of Dissertation	3
2	Literature Review	5
2.1	Security Vulnerabilities and the Importance of Updates in the Tor Network	5
2.2	User Attitudes Towards Updates	7
2.3	Automatic Updates and Improved Security	9
2.4	Research Gap and Rationale	10
3	Research Methodology	12
3.1	Research Strategy	12
3.2	Data Collection	13
3.2.1	Sampling Method	13
3.2.2	Data Collection Technique	13
3.3	Framework for Data Analysis	16
3.3.1	For quantitative part	16
3.3.2	Qualitative Analysis	18
3.4	Methodology Limitations	19
4	Survey Findings: Description and Analysis	20
4.1	Participants	20
4.2	Quantitative Analysis	21
4.2.1	User Experience: Factor Analysis	21

4.2.2	Upcoming Update Features: Factor Analysis - Automatic vs Current Update	23
4.2.3	Further Attitudes Towards Automatic Update Mechanism	24
4.3	Qualitative Analysis	25
4.3.1	Insights from Outdated Relay Operators	25
4.3.2	Analyzing Feedback on Current Updates	26
4.3.3	Analyzing Feedback on Automatic Updates	29
5	Discussion	32
5.1	Security Considerations	32
5.2	Reduce Perceived Cost	33
5.3	Balance Relay Operator Control	33
5.4	Improve Update Communication	34
5.5	Platform Inconsistency Need to be Addressed	35
5.6	Limitations	35
6	Conclusion	37
6.1	Conclusion	37
6.1.1	Objective 1: Update Drivers and Barriers	37
6.1.2	Objective 2: Tor relay Views and Practices	38
6.1.3	Objective 3: Manual Update versus Automatic Update	38
6.1.4	Objective 4: Recommendations for Automatic Update Design	39
6.2	Future work	39
	Bibliography	41
	A Combined Participants' Information Sheet and Consent Form	51
	B Questionnaire	56
	C Detailed Demographic Information	65
	D Wilcoxon Signed-Rank Test Results	67
	E Themes and Coded Responses	69
E.1	Themes and Coded Responses of Q10	69
E.1.1	Lacking Update Communication	69
E.1.2	Complex Update Process	70

E.1.3	Tor Auto-Update is Simple	70
E.1.4	Update Side Effects	70
E.1.5	Update as Instructed by Tor	71
E.2	Themes and Coded Responses of Q15	71
E.2.1	Support for New Auto-Updates	71
E.2.2	Desire to Opt-Out of Auto-Updates	71
E.2.3	Recommendations for Implementation	72
E.2.4	Opposition to the New Auto-Updates	73

Chapter 1

Introduction

1.1 Motivation

1.1.1 Tor and Relay Operators

This paragraph is taken verbatim from my IPP report: The Tor (The Onion Router) network is a widely used anonymization tool that allows users to protect their online identity and browsing activity from prying eyes. Tor works by routing internet traffic through a series of relays before reaching its final destination [81]. This process makes it difficult for an attacker to trace the user's activity back to their actual IP address. In this network, relay operators hold a critical role as they voluntarily operate the relays and take sole responsibility for relay updates. Due to its sensitive nature, Tor relays are suggested to be kept updated at all times. When certain relays continue to use outdated versions, they may be vulnerable to serious security flaws that compromise user privacy and anonymity. It is worth noting that Tor already implements measures to enhance network security, such as implementing End-Of-Life (EOL) for certain relays. However, creating the EOL within Tor involves thoughtful analysis of its potential effects on network traffic. Additionally, determining when to phase out specific relay versions presents its own set of challenges. Despite these efforts, little research has been conducted to understand the underlying reasons for some relay operators' reluctance to update. If relay operators hold positive attitudes towards updates, the need for implementing EOL for relays may be reduced [24]. Moreover, exploring the factors influencing relay operators' update decisions could reveal key attributes of automatic update mechanisms that can improve the update rate within the network [23, 24, 51, 55, 83].

Motivated by these considerations, this research aims to delve into the attitudes of

relay operators towards automatic update mechanisms in Tor. By gaining insights into their perspectives, concerns, and preferences, this study seeks to propose suggestions for future update designs in Tor.

1.1.2 Automatic Relay Update Framework

To contextualize and gain more specific insights into participants' attitudes towards automatic updates, this research particularly focuses on surveying the automatic relay update framework developed by Rochet and Elahi [73]. They proposed FAN for Flexible Anonymous Network, "a new software architecture for volunteer-based distributed networks that shifts the dependence away from protocol tolerance without losing the ability for the developers to ensure the continuous evolution of their software" [73]. This new update design addresses the complex task of maintaining a distributed network involving multiple actors, such as Tor. In the current update scenario within Tor, network developers lack control, leading to heterogeneity in software versions and the need for protocol tolerance and forward-compatible strategies. To address these issues, FAN provides an architecture that is independent of OS distributions and unattended relays within the Tor network. This design enables FAN to improve anonymous communication in a more flexible way. It achieves this by facilitating on-the-fly negotiation and deployment of protocol features.

In conclusion, the FAN design greatly improves the flexibility of relay updates in the Tor network. It gives developers more control over the update process, which strengthens security. However, to fully understand what makes automatic updates work well and be accepted by relay operators, it is essential to analyze operator attitudes about this design. It will shed light on the complex interactions between update technology features and operators' perceptions.

1.2 Objectives

The primary goal of this research is to enhance the understanding of relay operators' attitudes concerning automatic updates. With a specific focus on Tor relay operators' update practices, the study aims to reveal the factors influencing relay updates. Furthermore, it seeks to uncover the considerations underlying their decisions to accept or reject an automatic update framework.

Specifically, within the context of Tor, the objectives of this research are to:

1. *Identify* established factors impacting user attitudes towards updates.
2. *Explore* Tor relay operators' views and practices related to updating relays.
3. *Examine* how relay operators' attitudes differ or change in the context of manual and automatic updates.
4. *Formulate recommendations* on how automatic updates should be designed.

Objective 1 will be tackled in the Literature Review. Objectives 2 and 3 will be fulfilled by collecting and analyzing empirical data. Finally, objective 4 will be derived based on the findings from objectives 1, 2 and 3.

1.3 Results and Outcomes

In my study, I surveyed 54 Tor relay operators to understand their beliefs about updates and explore if these can be categorized into actionable factors for designing better update frameworks. I found that various factors, especially security concerns, influence how relay operators view updates. Interestingly, automatic updates raised more security concerns than the current Tor update method. Users shared similar worries as those found in previous research studies [22, 50, 83, 87]. Additionally, some factors that are significant in other contexts held less importance for Tor relay operators in this study. For example, the factor "Update seems unnecessary because everything works fine" [82, 83, 84] was of minimal concern to participants.

Furthermore, my research identified new influential factors that haven't been explored before, particularly relevant to Tor relays. One such factor is the relay operators' desire for control. Participants expressed the need for substantial control over the update process, especially within an automatic update system. Going beyond technical aspects, understanding user concerns and practices related to updates is crucial [24]. Through this survey, my study addressed this gap in the existing literature within the specific context of Tor.

1.4 Structure of Dissertation

The upcoming chapter delves into the existing body of literature concerning user attitudes towards updates. This exploration enables readers to trace the evolution of

research in this domain, spanning diverse contexts such as Windows, mobile applications, and IoT devices. Analyzing these prior studies will offer insights into the survey's development and emphasize the novelty of this research in the context of Tor.

Chapter 3 outlines the chosen research methodology, encompassing the selected research strategy, data collection methods, and analytical techniques. The chapter also addresses any limitations inherent in the chosen methodology, providing transparency in the approach and its potential implications.

In the fourth chapter, the collected data and its corresponding analyses are presented and compared. This section gives a full picture of the attitudes relay operators have. It presents the data clearly to facilitate observations.

Chapter 5 engages in a comprehensive discussion based on the gathered data, organizing the results into separate categories that offer diverse perspectives. These categories cover crucial aspects such as security, cost, control, communication, and platform inconsistency.

The final chapter, Chapter 6, serves as the conclusion of this dissertation. It presents a brief overview of the research findings, discusses their significance, and suggests possible directions for future research.

Chapter 2

Literature Review

The Literature Review will focus on examining the main issues surrounding three key aspects: the significance of studying Tor network updates, the existing research on user attitudes toward updates, and identify the significance of automatic updates. The primary focus of this review will be on addressing Objective 1: Identify established factors impacting user attitudes towards updates.

By exploring the above areas of literature, a significant contribution will be made to this research. A key aspect of this exploration involves identifying the factors shaping users' update decisions. Additionally, this literature review aims to shed light on the benefits users can derive from automatic updates.

At the end of this major section, the reader will be better informed in these areas. This understanding will enable a clear focus and justification for conducting empirical research in the field of relay operators' attitudes toward automatic updates in the Tor network. A sensible starting point is to investigate the security vulnerabilities within the Tor network, thereby highlighting the crucial role of updates in mitigating these vulnerabilities and ensuring network security.

2.1 Security Vulnerabilities and the Importance of Updates in the Tor Network

The Tor network's distinct architecture [18] has prompted extensive security research, encompassing various aspects [40]. Notably, studies have demonstrated vulnerabilities in the network, with attacks like traffic analysis and end-to-end correlation posing risks to user anonymity [7, 8, 39, 56]. Furthermore, the Exit Node Vulnerability has been

highlighted as a potential deanonymization risk [27]. In the face of evolving threats, even deep learning approaches have been leveraged to create potential attack models [57]. Addressing these challenges necessitates ongoing updates.

Thus, timely and effective updates are critical for Tor security. It is widely acknowledged in the security domain that keeping systems up to date is vital to preventing attacks [14, 1, 4, 48, 58, 80]. Tor developers diligently address identified vulnerabilities and promptly issue updates [85]. However, the emphasis is on “timeliness” [9, 47, 86], a responsibility that rests upon users, Tor relay operators in this context. Furthermore, updates can sometimes introduce risks, potentially creating new vulnerabilities or data leaks [89], thereby underscoring the vital role of user engagement.

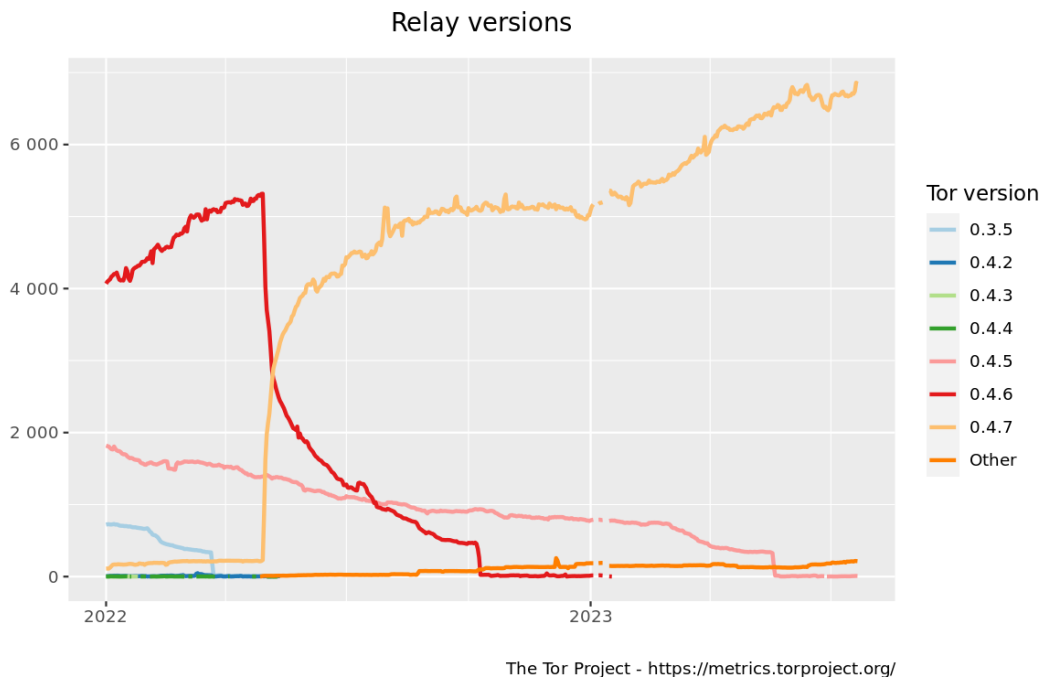


Figure 2.1: Relay versions running on Tor from January 1, 2022, to June 23, 2023

Despite the persistent emphasis on timely updates, achieving sufficient update rates remains an ongoing challenge across software [2, 42], including for the Tor network. As Figure 2.1 shows, new Tor versions see delays before broad deployment, with releases like 0.4.7 requiring nearly six months for substantial updating by June 2022. This delayed response to updates poses security risks and undermines efforts to mitigate vulnerabilities promptly.

In addition to adoption delays, the persistence of outdated Tor versions also presents concerns. While protocol tolerance [64] allows multiple versions of relays to coexist, those outdated relays create vulnerabilities [74]. Thus, to ensure network security, Tor

has to implement End-of-Life (EOL) status for outdated relay versions [61]. However, this action may result in the deactivation of a percentage of working relays, which is not ideal for the overall network performance. For instance, in June 2023, Tor removed active 0.4.5.x relays, affecting public relays on this version. [30]

Given the significance of updates in maintaining the security of the Tor network, this study aims to examine the delay in updates by relay operators and explore their perceptions of automatic updates. The objective is to provide recommendations for an automatic update process that balances their preferences with network security. To accomplish this, the upcoming literature review section will investigate prior research on user attitudes towards updates across different contexts, providing a foundation for this study.

2.2 User Attitudes Towards Updates

As highlighted in the preceding section, timely relay updates are crucial for vulnerability mitigation, with their efficacy relying on relay operators' attitudes and behaviours. Given the lack of literature explicitly examining Tor relay operator perspectives, this section delves into investigating updating attitudes in various contexts. The aim is to identify transferable practices aligned with the research objectives.

In exploring user behaviours in software, early pioneers [16, 25] conducted research in the Android app. Though their focus was on installation rather than updates, they unveiled a crucial user attitude factor: the challenge users face in comprehending presented messages, highlighting inefficiencies in communication. Expanding on this finding, Kelley et al. [41] affirmed the role of effective communication in influencing user installation behaviours. Their developed framework, designed for enhanced communication, proved successful in aiding users during the installation process. Furthermore, Möller et al. [53] emphasized the security consequences of low update rates. Their data highlighted instances where users struggled to interpret provided information accurately, further reinforcing the necessity for effective communication to ensure clear user understanding.

In 2013, Nguyen et al. [59] explored mobile apps and introduced a new factor, crowdsourcing, that shapes user behaviour. Their findings highlighted how users seek solutions from others when facing issues, contributing positively to security. Similarly, Rader et al. [69] conducted a survey focused on security stories, yielding similar outcomes. Their research methodologies have provided inspiration for subsequent

studies, including the approach undertaken in this research. These works suggest that Tor relay operators might also turn to community forums for updates and security information.

As the significance of user attitudes became evident, efforts expanded beyond mobile apps to include Windows environments as well. For instance, similar to the findings in Android mobile apps, user confusion emerges as a crucial attitude towards updates in Windows [87]. To address the confusion aspect of user attitudes, other researchers have focused on designing effective update messages. Others [26, 33, 34, 35, 82] have explored methods to improve user understanding and reduce confusion during the update process. The work of Fagan et al. [23] is remarkable for its comprehensive survey covering 20 types of software update messages. Their findings highlighted the strong influence of emotional factors, varying by software type. This underscores the need for targeted research within the Tor community, considering its distinct characteristics and user base, rather than directly generalizing findings from other software studies.

The above studies emphasize addressing user confusion in update messaging, yet a comprehensive understanding of user attitudes toward updating requires broader considerations beyond notifications. In their study, Vaniea et al. [84] interviewed 37 non-expert Windows 7 users. Their findings align with prior research, highlighting the significance of crowdsourcing and the impact of past experiences on present behaviours. The researchers categorized user attitudes into three themes: surprise at new features, unclear update purposes, and perceived update necessity. This analysis method benefits future researchers, including me, in recognizing that user attitudes are shaped by various interconnected factors. Another extensive exploration of updating attitudes is undertaken by Vaniea and Rashidi [83]. They surveyed 307 respondents to explore the complete updating process from users' perspectives, identifying issues at each stage. Their qualitative data analysis methodology influences my research analysis approach. Furthermore, their emphasis on factors like limited computer resources leading to updating issues aligns with my study's considerations.

Studies examining user attitudes towards automatic updates also exist. Mathur and Chetty [49] concentrated on mobile applications. Their study, focusing on automatic updates without any user intervention, highlighted the impact of security considerations on update behaviours, which aligns with my research's exploration of security aspects.

A highly influential contribution in this field is Mathur and Chetty's follow-up work that quantifies U.S. users' beliefs on software updating [51]. Despite potentially limited generalizability due to sample size, their study yields valuable insights. Notably,

users align with three key axes of concerns: update necessity, costs, and risks. They also define 15 categories of influencing factors, offering a structured foundation for my research to extend and tailor insights to the distinctive context of the Tor network. Other researchers have further investigated specific factors highlighted by Mathur et al., including “the need to restart” [55] and “update costs” [71]. Through a laboratory experiment, Rajivan et al. [71] identified new aspects like “willingness to take risks,” tied to user characteristics, driving more timely updates. While my research may not encompass this particular factor, these investigations highlight that delving into established factors can unveil novel elements influencing user decision-making.

User attitude research expanded to include smart home IoT updates, with findings diverging from other contexts like Windows and mobile apps [6]. Smart home users exhibited greater awareness of update importance and urgency for their devices. Furthermore, the links between update perceptions and views on security/privacy proved less straightforward. These revelations further underscore the influence of contextual factors shaping user attitudes. They highlight the need to study Tor relay operator perspectives recognizing their unique context.

2.3 Automatic Updates and Improved Security

Before delving into surveying relay operators’ attitudes towards automatic updates, it’s crucial to establish that automatic updates hold a promising future. Duebendorfer and Frei [19] demonstrated that automatic updates combined with minimal operating system reliance proved most effective for keeping users updated. Similarly, Zhen-hai and Yong-zhi [88] explored streamlining frameworks to enable efficient automatic updates with minimal manual intervention. Empirical experiments evidenced the value of such automated approaches for enhancing software security. However, it is essential to strike a balance between automation and user control. Completely taking control away from users and implementing a “Fixed Policy” approach, may not be the best way to address user preferences [22]. Despite this, automation in updates undeniably boosts update rates and subsequently enhances security.

The automatic update framework surveyed in the study is designed to enhance security within the Tor network. It [73] involves updates being automatically pushed after developers consider relay operators’ opinions. Users maintain control over update appropriateness, reflecting their preferences through votes. However, once adopted, relay operators cannot stop update implementation. This positions the automatic update

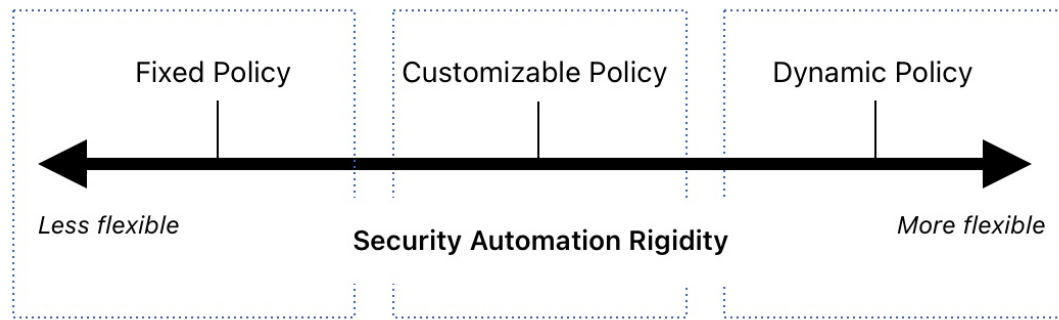


Figure 2.2: The spectrum of automation approaches [22]: “Fixed Policy” involves automatic installation without user intervention; “Customizable Policy” permits security policy customization, often based on system-wide defaults set by administrators [31]; “Dynamic Policy” enables personalized security tailoring [75].

between the “Fixed Policy” and “Customizable Policy” (Figure 2.2). Furthermore, the evaluated framework aligns with effective design principles, prioritizing appropriate automation. Prior research [22, 87] highlights the importance of user involvement and comprehensive information provision. Excessive automation often leads to user misunderstanding. The Tor relay automatic update framework includes user participation and ensures users are provided with comprehensive information, mitigating such misunderstandings.

In summary, the literature underscores the potential of automatic updates to enhance software security through user involvement, transparent information, and empowered decision-making. The surveyed Tor network’s automatic update design [73] adheres to these principles, bolstering the Tor ecosystem’s security. The remaining aspect is understanding relay operators’ perceptions and specific requirements towards automatic updates.

2.4 Research Gap and Rationale

The literature review offers a comprehensive understanding of timely updates within the Tor network, user attitudes toward software updates, and the value of automatic updates. It highlights the need for up-to-date relays and comprehending user decision factors. Additionally, it underscores the potential advantages of adopting automatic update mechanisms for security and user experiences.

A notable gap identified in the literature is the limited focus on the attitudes of Tor

relay operators towards automatic updates. While studies have explored user perspectives in various contexts, there is a lack of in-depth research specifically addressing the unique considerations within the Tor network. Given the critical role Tor plays in safeguarding privacy and security, understanding relay operators' attitudes towards updates becomes paramount in developing effective and tailored automatic update solutions.

The necessity of this empirical research stems from the need to bridge the aforementioned research gap and extend the current understanding of user attitudes towards automatic updates. By focusing on Tor relay operators, we can gain insights into their specific concerns, preferences, and experiences in the context of relay updates. Such knowledge is essential for designing automatic update mechanisms that strike the right balance between user control and automation, ensuring efficient and secure updates within the Tor network.

The subsequent section of this dissertation will outline the Research Methodology employed to gather empirical data. It covers the chosen research strategy, data collection methods, sample selection, and data analysis techniques.

Chapter 3

Research Methodology

Chapter 2 provided an in-depth literature review on user attitudes and automation in updates, addressing Objective 1. This review highlighted the research gap in understanding relay operators' attitudes toward updates in Tor.

Objectives 2 and 3 advance the research by involving Tor relay operators in data collection and analysis. This empirical research serves to bridge theory and practice, creating a more comprehensive understanding of challenges in implementing automatic updates within the Tor network.

In the Research Methodology chapter, the details of the research strategy adopted to address the identified research issues will be provided. This will include the means of collecting data for analysis, such as tools and sample selection, as well as the chosen analysis approach. Additionally, the chapter will address potential limitations associated with the research strategy.

3.1 Research Strategy

The chosen research strategy for this empirical study is *Survey-based research*, which involves collecting information from a sample of individuals through their responses to questions [15]. This approach is suitable for this project for several reasons. First, given its adaptability to both quantitative and qualitative research methods, it fits the mixed-methods approach of the project. Moreover, surveys can efficiently collect data from a diverse participant pool, capturing varied opinions. Leveraging online survey platforms, researchers can reach a significant relay operator population without requiring face-to-face interactions. This approach ensures participant anonymity while promoting open expression of opinions. In addition to efficiency, surveys offer a structured approach to

data collection, ensuring consistency in the types of responses obtained. By employing both quantitative and qualitative survey items, the research can gain deeper insights into the attitudes of relay operators while also quantifying their responses to understand the significance of certain factors in the Tor network. Prominent studies in this domain also employ this research strategy [23, 24, 51, 82, 83], further affirming its appropriateness.

3.2 Data Collection

The empirical research aims to comprehensively understand relay operators' attitudes toward automatic updates within the Tor network. This involves examining the factors contributing to their decisions to update and assessing their reactions to the new automatic update design. To achieve this goal, a mixed-methods approach is used for data collection, integrating qualitative and quantitative components.

3.2.1 Sampling Method

Convenience sampling [21] was chosen to select Tor relay operators as participants. This approach was selected because the researcher was actively engaged in the Tor project forum, providing convenient access to potential participants. Thus, emails were sent to the Tor relay operators' mailing list, following the suggestion of the Tor project leader. This ensured a targeted and respectful approach to reaching the intended population. Convenience sampling was also chosen considering time constraints. Given the limited time available for data collection, it allows the research to proceed in a timely manner while ensuring that sufficient insights are obtained to address the research objectives. It is important to acknowledge that convenience sampling does not involve random selection, and therefore, the findings cannot be generalized to represent the entire population of Tor relay operators. Instead, the primary objective of this research is to gain exploratory insights into relay operators' attitudes towards automatic relay update issues.

3.2.2 Data Collection Technique

The survey-based research in this project utilized a structured questionnaire as the primary data collection technique.

3.2.2.1 Questionnaire Design

Formulated with a focus on clarity and precision, the questionnaire serves as a strategic tool for gathering responses from relay operators. To ensure the questionnaire's effectiveness, it underwent several rounds of reviews involving both the supervisor and Tor Project leaders. This section delves into the structure of the survey, highlighting its potential to yield valuable insights.

The questionnaire began with "Demographic Questions" (Q1-7), aimed at gathering relay operators' information while ensuring their anonymity. These questions avoided sensitive data, like age or gender, to prevent identity disclosure. Instead, they focused on essential aspects such as relay operation duration and relay types. This background and experience insight was crucial, as prior studies showed expertise significantly shapes attitudes [12, 46, 36, 62]. Notably, when participants indicated an outdated relay, an open-ended question (Q7) inquired about their reasons. This approach captured nuanced insights into their specific challenges and concerns, facilitating a deeper understanding of their update attitudes.

The questionnaire's second section centred on "Attitudes towards the current update" (Q8-10). To ensure participant information consistency, a neutral introduction to Tor relay update functioning was provided. This was particularly crucial due to participants' diverse backgrounds and varying knowledge levels about relay updates. Ensuring uniform information dissemination promoted reliable responses. Within this section, participants were presented with a range of factors identified in previous research and asked to rate the importance of each factor in influencing their update decisions. Questions were divided into two sets: Q8 focused on potential user experience impacts, while Q9 delved into update feature changes. The specific factors examined in Q8 are as follows: Factors UP1 to UP5 (UP - Update) were identified in research conducted by Mathur et al. [51]. Given that updates in the Tor network often come with descriptions of the changes introduced, especially regarding security enhancements, UP6 and UP7 seek to assess participants' perceptions regarding the relevance of security and new features associated with updates.

- UP1-Necessity: Everything works fine, so this new update seems unnecessary [82, 83, 84]
- UP2-Time: Update takes up a significant amount of the user's time during the installation process [83]

- UP3-Restart: Update requires unnecessary restarts [83, 87]
- UP4-Infrequent: I do not keep up with Tor update news regularly
- UP5-Unimportant: Update seems unimportant [84]
- UP6-Security-Related: Update is related to security
- UP7-Feature-Related: Update is related to features

The specific factors examined in Q9 are as follows: Factors UF1 to UF6 (UF-Update Feature) were identified in research conducted by Mathur et al. [51].

- UF1-Unwanted: Updates may add unwanted features [83]
- UF2-Wanted: Updates may remove wanted features [83]
- UF3-Compatibility: Updates may cause compatibility issues [50, 83]
- UF4-CPU: Update process may consume too much CPU [50, 83]
- UF5-Bugs: Updates may introduce new bugs [50, 82]
- UF6-Malicious: Updates may introduce malicious content [50, 83]

Both sets used a five-Likert scale [45]. This scale captured subtle opinion differences, determining pivotal factors in update decisions. Furthermore, Q10 allowed participants to contribute suggestions and share negative experiences with the current update framework. This open-ended question aimed to collect qualitative insights for identifying new attitude factors specific to Tor.

The third section of the questionnaire delved into “Perceptions of Automatic Update Design” (Q11-15). An introduction to the FAN (Flexible Anonymous Network) [73] was provided at the beginning of the section. The introduction outlined the components involved in the design and each component’s responsibilities, including those of Tor relay operators. This detailed introduction aimed to encourage insightful responses, particularly from participants with expertise in the field. Q11 was derived from section 2, specifically Q9. Comparing Q9 and Q11 responses enabled the identification of variations resulting from the new automatic update design introduction. Q12-14 revolved around the FAN design, offering contexts for participants to assess. These questions aimed to ensure participants’ thorough consideration of the FAN framework’s

implications. The final question, Q15, was open-ended. It aimed to gather diverse responses, potentially unveiling new factors participants believed important for designing automatic update frameworks.

The questionnaire for the Tor relay operators can be found in Appendix B.

3.2.2.2 Data Collection Platform

CryptPad [79] was chosen as the survey platform for this research project based on a recommendation from the Tor project leader during a meeting to discuss research objectives. It presents distinct advantages, as stated below. One notable advantage is that CryptPad allows participants to complete the survey without an account, ensuring their anonymity throughout the process. This feature provided reassurance to participants that their identity will remain undisclosed. Moreover, CryptPad prioritizes data security and confidentiality, which aligns with the critical requirement of Tor relay operators. The platform utilizes encryption [79] in the browser to prevent the service provider from accessing the content of the survey responses and other documents, ensuring that participant data remains private and protected. Another advantage of using CryptPad is the capability to directly attach a PDF version of the participant consent form at the beginning of the questionnaire. This seamless integration enables participants to provide their consent without any interruptions in the survey process. Overall, CryptPad's ability to provide an anonymous and secure survey environment, combined with its practical features, makes it a well-suited platform for conducting this research with Tor relay operators.

3.3 Framework for Data Analysis

Figure 3.1 outlines the data analysis framework employed in this research. Subsequent sections will offer an in-depth introduction and explanation of each phase within this process.

3.3.1 For quantitative part

The quantitative analysis of the questionnaire employed descriptive statistics to analyse participants' attitudes and behaviours. This analysis was structured into three primary sections, starting with the demographic section.

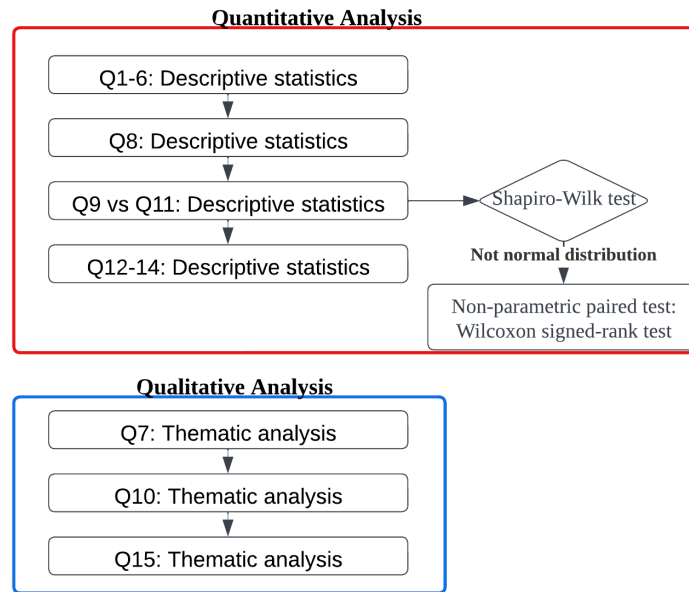


Figure 3.1: Data Analysis Process Overview

Descriptive statistics were utilized to assess the update status of participants. This exploration contributed to a clearer understanding of the update behaviours within the sample population and, consequently, shed light on the broader update statuses within the Tor network user base.

Descriptive statistics were also applied to the close-ended questions in the second part of the survey. Specifically, for Question 8 of the questionnaire, which encompasses seven factors, key statistical measures were identified. These measures included the minimum, maximum, mean, and skewness of the responses. Through seven individual analyses, a comprehensive evaluation of belief distribution among participants was achieved. In addition, the five-point Likert scale in this study was regarded as an interval scale. This approach has been extensively utilized in the field of psychology, as discussed by several scholarly works [3, 11, 78]. Scores falling within the range of 1 to 1.8 indicate the factor is perceived as “not at all important.” Values in the range of 1.81 to 2.6 suggest a general sentiment of “low importance.” Scores ranging from 2.61 to 3.4 signify a “neutral” stance regarding the factor’s importance. Values between 3.41 to 4.2 indicate “important”. Finally, scores ranging from 4.21 to 5 are perceived as “very important”[63].

Similar to Part 2, the third segment of the survey underwent descriptive statistical analysis for Questions 12-14. The central statistical measure employed for comparison

was the mean value for Question 11. This analysis aimed to identify the factors that participants considered more significant in shaping their update decisions within the new automatic update mechanism.

However, the analysis extended beyond this point. To assess the impact of the new automatic update scheme introduced in the survey, the results from Q9 (close-ended questions about the current update process) and Q11 (close-ended questions about the proposed automatic update process) were compared. The data analysis was conducted using the Statistical Package for the Social Sciences (SPSS). Paired tests were applied to discern significant variations in participants' attitudes towards the same given factors between the two update processes. This analysis offered insights into whether the new automatic update scheme influenced participants' perspectives on specific factors. The selection of the appropriate paired test method was contingent on whether the collected data adhered to a normal distribution, as determined by the Shapiro-Wilk test result [77]. As the collected data didn't demonstrate normal distribution, a non-parametric paired test called the Wilcoxon signed-rank test was utilized [52, 72].

3.3.2 Qualitative Analysis

For the qualitative part of the questionnaire, the thematic analysis [44] was employed. This analysis was conducted using NVivo and facilitated the exploration of new factors influencing relay operators' perspectives on updates. Importantly, the six phases outlined in Figure 3.2 demonstrate interactive attributes. The adopted approach in this study was inductive, where themes emerge in the course of analysis [13].

For Q7 in Part 1, a thematic analysis approach was employed to analyze responses and identify recurring themes. This allowed for the exploration of factors possibly contributing to relay operators' reluctance in updating their Tor relays. Notably, this analysis occurred prior to the presentation of specific factors, ensuring that the reasons participants provided were derived from their own update experiences.

For Q10 in Part, participants were asked about their additional views on the current update mechanism for Tor relays. A similar analysis was employed, and potential novel insights that could impact relay operators' decisions on updates were sought.

Moving on to Q15 in Part 3, participants were asked about their additional perceptions towards the new automatic update mechanism. Employing thematic analysis, the themes inherent in participants' responses were examined, offering the potential to unveil distinctive factors linked to the new automatic update mechanism.

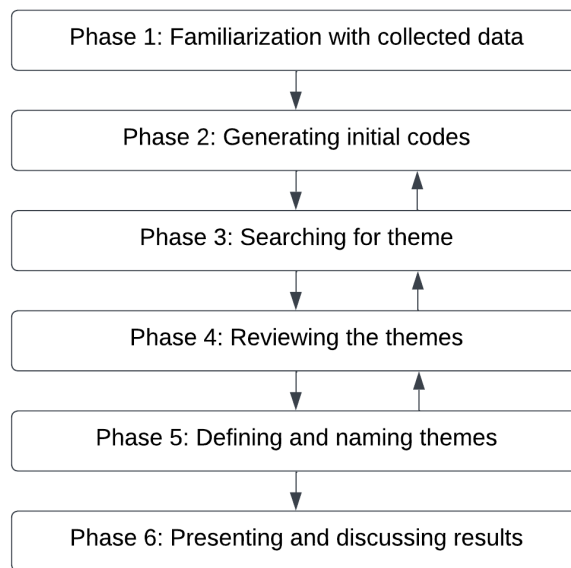


Figure 3.2: Thematic analysis: Six interactive phases [44]

3.4 Methodology Limitations

The research methodology in this study has limitations, particularly concerning the use of questionnaires, which may not fully capture participants' perceptions. Furthermore, the lack of direct interaction with the automatic update framework could pose challenges for participants in understanding the topic. While efforts were made to design a clear questionnaire, the absence of such direct interaction may have limited participants' depth of responses. Conducting a field study with real-time interactions might have offered a more immersive and informative experience. In addition, self-reporting bias remains a potential concern, as participants might provide socially desirable responses, skewing the data to some extent. Additionally, the lack of multiple coders in the thematic analysis process may have impacted intercoder reliability, as having multiple coders would have increased confidence in the analysis.

Despite these limitations, the research methodology employed in this study has yielded valuable insights into the attitudes and perceptions of Tor relay operators regarding the update process and the new automatic update mechanism. Acknowledging these limitations can serve as a basis for further research.

Chapter 4

Survey Findings: Description and Analysis

This chapter presents the findings of the survey outlined in Chapter 3, focusing on Objectives 2 and 3: *Explore* Tor relay operators' views and practices related to updating relays and *Examine* how relay operators' attitudes differ or change in the context of manual and automatic updates.

4.1 Participants

A total of 55 participants were recruited. Following data cleaning, one response was eliminated due to duplicate entries, resulting in a dataset comprising 54 responses. Among the 54 respondents, a significant variation was evident in their relay operating experience, highlighting diverse levels of involvement. Table C.1 in Appendix C summarizes the key characteristics of the participants.

The update status of their relays as shown in Table 4.1 reflects a positive trend towards updates. This high update rate can be explained by several factors. Firstly, the availability of recent updates over an extended period gives relay operators enough time for necessary updates, supported by Tor version metrics (Figure 2.1). Moreover, the recent implementation of End-of-Life (EOL) for 0.4.5.x within Tor [30] enforces updates to prevent relay take-downs. Additionally, survey participants actively engage in the Tor community, leading to a stronger motivation to maintain the latest relay versions due to heightened security concerns [24, 51]. However, the data shows a small number of participants with delayed updates. Reasons for not updating will be analyzed in the later section.

Table 4.1: Relay Version Status of Tor Relay Operators Participants

Relay Version Status	# Participants
Up to date	49
One version behind	3
Three or more versions behind	1
Not Sure	1

4.2 Quantitative Analysis

4.2.1 User Experience: Factor Analysis

To visualize beliefs for each factor in Question 8, a 100% stacked bar graph was created, displaying respondent percentages for each belief category (Figure 4.1). This graph provides insights into diverse participant attitudes, sorted by the “very important” percentage. A notable finding is that a significant proportion of participants (77.8%) designated the factor [UP6] as “very important,” underscoring relay operators’ strong focus on security implications tied to updates. This could be attributed to Tor’s primary purpose of ensuring anonymity [76], with relay operators valuing its security aspect. In contrast, none of the respondents rated the factor [UP1] as “very important,” suggesting a consistent participant perception that updating relays remains crucial, regardless of their current operational status.

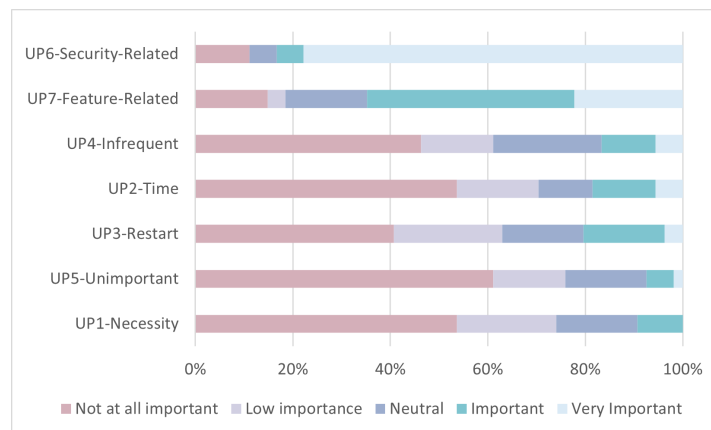


Figure 4.1: The distribution of software beliefs from Tor relay operators samples: User Experience Aspect

The descriptive statistics of Question 8 (Table 4.2) provide more detailed statistical support for the importance of each factor. Two factors stand out from this analysis:

Table 4.2: Influential Attitudes toward Updating: User Experience

Attitude Factor	Metrics			
	Min	Max	Mean	Skewness
[UP6]Update is related to security	1	5	4.39	-2.03
[UP7]Update is related to features	1	5	3.54	-0.89
[UP3]Update requires unnecessary restarts	1	5	2.20	0.62
[UP4]I do not keep up with Tor update news regularly	1	5	2.15	0.72
[UP2]Update takes up a significant amount of the user's time during the installation process	1	5	2	1.01
[UP1]Everything works fine, so this new update seems unnecessary	1	4	1.81	0.93
[UP5]Update seems unimportant	1	5	1.72	1.29

Note. 5 Very important, 4 important, 3 Neutral, 2 Low importance, 1 Not at all important.

[UP6] and [UP7]. Both have relatively high mean values, with negative skewness indicating responses leaning right of the mean. Overall, [UP6] is perceived as “very important,” while [UP7] as “important”. In contrast, [UP5] had the lowest mean of 1.72. This indicates that many participants don't consider “update necessity” to be a crucial factor in shaping their decisions to update.

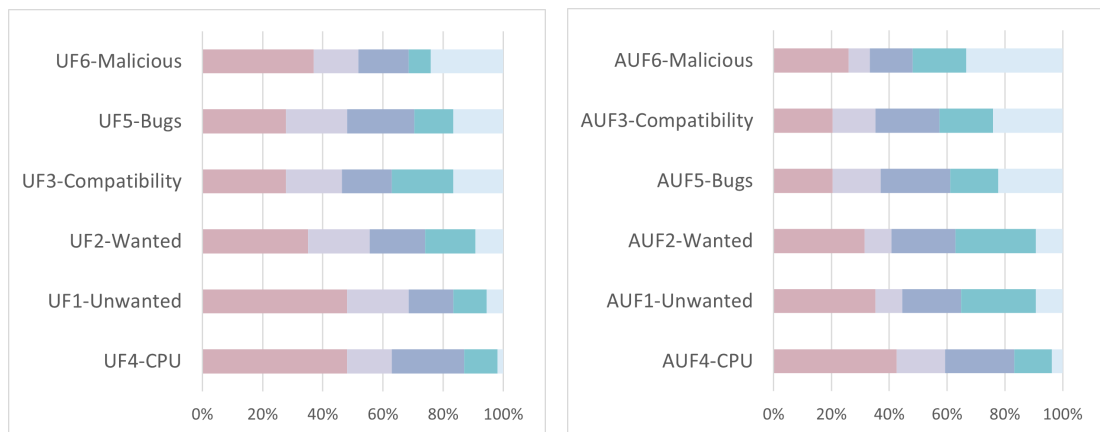
This contrasts with previous research [51], and detailed reasons are explored in the later qualitative analysis.

Overall, the results from Question 8 contribute to Objective 2: *Participants prioritize the potential benefits of new updates over personal experience factors. Notably, the most significant factor is security, aligning with earlier research [24, 51, 83]. This highlights a trend toward objective decision-making in the Tor community.*

4.2.2 Upcoming Update Features: Factor Analysis - Automatic vs Current Update

4.2.2.1 Initial Analysis: Comparing Participant Beliefs between Current and Automatic Update

After analyzing Questions 9 and 11, Figure 4.2 offers an overview of response patterns for current updates (UF) and automatic updates (AUF). It is clear from the graphs that responses span the entire five-Likert scale. Therefore, metrics for minimum and maximum were omitted from the following table.



(a) Distribution of attitude importance levels under current update (UF).

(b) Distribution of attitude importance levels under automatic update (AUF).

Figure 4.2: Distribution of respondents rating level of importance for the six attitude factors.

Table 4.3 presents descriptive statistics results. [AUF/UF6] and [AUF/UF5] remain significant factors for both mechanisms, aligning with Question 8. This contributes to Objective 2: *security-related concerns are central to relay operators' update considerations*. Mean value comparison suggests factors were more important in the automatic update scenario. However, for research Objective 3, further exploration is needed to understand the extent of this change, as indicated in subsequent paired tests.

4.2.2.2 Wilcoxon Signed-Rank Test: Attitudes towards current vs. automatic update Framework

The Shapiro-Wilk test yielded p-values below 0.001 for the six key attitude factors in both the current update (UF) and automatic update (AUF) contexts. This signifies

Table 4.3: Mean for Ratings of Influential Attitude Factors. It compares current and new automatic update mechanisms across six factors.

Attitude Factor	Mean	
	Current Update	Automatic Update
[AUF/UF1](Automatic)Updates may add unwanted features	2.06	2.65
[AUF/UF2](Automatic)Updates may remove wanted features	2.44	2.74
[AUF/UF3](Automatic)Updates may cause compatibility issues	2.80	3.11
[AUF/UF4](Automatic)Update process may consume too much CPU	2.04	2.19
[AUF/UF5](Automatic)Updates may introduce new bugs	2.70	3.04
[AUF/UF6](Automatic)Updates may introduce malicious content	2.67	3.26

a departure from normal distribution for all factors. Consequently, non-parametric tests—specifically the Wilcoxon signed-rank tests—were used to analyze variable differences. The corresponding results can be found in Table D.1 within Appendix D.

Table 4.4 shows p-values from Wilcoxon signed-rank tests [17] for influential attitude factors between UF and AUF. Factors 1, 2, 5, and 6 exhibit statistically significant differences ($p < 0.05$), while factors 3 and 4 exhibit no apparent difference ($p \geq 0.05$). Additionally, the results suggest an increase in considerations for factors 1, 2, 5, and 6 in the after-measurements. For detailed results, refer to Figure D.1 in Appendix D.

The analysis of Questions 9 and 11 directly addresses research Objective 3: *It confirms that relay operators' attitudes vary in the context of manual and automatic updates. Notably, concerns regarding unwanted changes, new features, bugs, and malicious attacks hold greater importance with automatic updates. However, factors like compatibility and CPU show no significant difference.*

4.2.3 Further Attitudes Towards Automatic Update Mechanism

Results for Questions 12-14 offer additional insights into participants' attitudes towards the automatic update mechanism (Table 4.5). The mean rating for [AUN1] is 3.65,

indicating its perceived “important”. Similarly, [AUN2] falls within the “important” range. Regarding positive aspects of automatic updates, [AUP1] has a mean rating of 3.20, placing it in the “neutral” category. [AUP2] received a slightly higher rating of 3.24, also within the “neutral” range. All factors display a moderately negative skew, consistently rating towards the upper half of the five-point scale.

In summary, this further addresses Objective 2: *While user-friendliness factors into relay operators’ update decisions, security considerations remain paramount. Additionally, participants appear to have reservations about the new automatic update framework.*

4.3 Qualitative Analysis

4.3.1 Insights from Outdated Relay Operators

Analyzing the responses of the 4 participants with outdated relays provides valuable insights. Two participants highlighted the time-consuming aspect of manual updates, indicating that perceived effort and time requirements significantly hinder timely updates. This aligns with previously identified attitudes [23, 51, 83], though this factor wasn’t prioritized in the quantitative analysis. These responses contribute to addressing Objective 2, *emphasizing the importance of perceived costs like time in affecting relay operators’ update behaviour.*

Two new factors were discovered from the responses. One participant highlighted,

Table 4.4: Wilcoxon Signed-Rank Test Results Comparing Influential Attitude Factors Between Current and Automatic Update Framework.

Factor Median Diff. (Auto-Current)	Wilcoxon Signed Ranks Test	
	Z	Asymp. Sig. (2-tailed)
AUF1-Unwanted - UF1-Unwanted	-3.014b	0.003
AUF2-Wanted - UF2-Wanted	-2.080b	0.038
AUF3-Compatibility - UF3-Compatibility	-1.721b	0.085
AUF4-CPU - UF4-CPU	-1.199b	0.230
AUF5-Bugs - UF5-Bugs	-2.255b	0.024
AUF6-Malicious - UF6-Malicious	-2.974b	0.003

Note. b. Based on negative ranks.

Table 4.5: Attitudes toward Automatic Update Mechanism: Specific Factors

Attitude Factor	Metrics			
	Min	Max	Mean	Skewness
[AUN1]Introduces a single point of failure	1	5	3.65	-0.69
[AUN2]Reduces relay operators' ability to customize relays	1	5	3.46	-0.40
[AUP1]Simplifies then update process, more user-friendly	1	5	3.20	-0.43
[AUP2]Incorporates feedback from relay operators	1	5	3.24	-0.56

“P5: It’s best practice in IT to be on the n-1 Version unless a critical issue is found.”

Another participant mentioned, *“P54: no more updates on debian repository deb.torproject.org.”* Given the diverse platforms used for Tor relays [43], the availability of updates may vary between platforms, presenting new challenges for relay operators in obtaining timely updates. These responses further address Objective 2 by *introducing novel factors not explored in prior studies: “n-1 Version Practice” and “Limited Repository Updates.”* While *“n-1 Version Practice” could be generalized to other contexts, “Limited Repository Updates” remains specific to Tor.*

4.3.2 Analyzing Feedback on Current Updates

The thematic analysis described in the previous chapter led me to identify eight sub-themes from Q10. These were grouped into five key themes: Lacking Update Communication, Complex Update Process, Tor Auto-Update is Simple, Update Side Effects and Update as Instructed by Tor. Each of these themes and subthemes is outlined in Table 4.6. Below, I have provided the results with quoted statements. These themes address Objective 2 by *explore specific perspectives influencing relay operators’ update choices.* The complete list of coded responses can be found in Appendix E.1.

4.3.2.1 Lacking Update Communication

The theme “Lacking Update Communication” encompasses two subthemes: “Insufficient Pre-Update Information” and “Inadequate Post-Update Reporting.” This theme highlights the barriers relay operators face in updating their Tor relays due to a lack of

Table 4.6: Subthemes informing each theme

Theme	Subthemes
Lacking Update Communication	<ul style="list-style-type: none"> •Insufficient Pre-Update Information •Inadequate Post-Update Reporting
Complex Update Process	<ul style="list-style-type: none"> •Difficulty Accessing Updates •Preference for Automation
Tor Auto-Update is Simple	
Update Side Effects	<ul style="list-style-type: none"> •Update Penalization Concerns •Machine Stability Concerns
Update as Instructed by Tor	

effective communication regarding updates.

The subtheme **“Insufficient Pre-Update Information”** revolves around relay operators not receiving adequate information prior to updates, leading to hesitation in timely updates. One participant specifically mentioned:

“P6: There isn’t a roadmap or a to-do list of incoming features/bug fixes. I don’t know what’s planned next or what’s coming next year for example.”

The second subtheme, **“Inadequate Post-Update Reporting,”** stresses the importance of clear communication after updates. Quick notifications reassure operators to respond promptly. Participants emphasized that,

“P39: Maybe, a changelog written in the user’s home directory who operates the tor-daemon (e.g. one user of group debian-tor under Debian) would be nice”

4.3.2.2 Complex Update Process

The theme “Complex Update Process” consists of two subthemes: “Difficulty Accessing Updates” and “Preference for Automation.” This theme highlights how complex update processes can discourage relay operators from updating.

“Difficulty Accessing Updates” reflects the challenges relay operators face in accessing the latest updates. What occurs most often in the response is platform inconsistency. Participants operating on different platforms, such as FreeBSD and Debian, reported delays in receiving updates. One participant highlighted:

“P17: ...Sometimes it takes a while before the FreeBSD packages are updated. Would be great to speed this up if possible (especially when there are security updates).”

Additionally, participants cited facing issues while seeking update-related informa-

tion, “P9: login to unopened sites”.

One potential reason for this is that the Tor project’s instructions for maintenance might not be well-organized. For instance, searching for automatic relay updates initially yields instructions for Debian and Ubuntu systems [66], which might confuse non-users of these systems. Another instance is the “Publishing Tor relay auto-update instructions” project [20], which refers to an outdated “Tor Relay Guide” [29] lacking auto-update details. The actual auto-update instructions were relocated to Tor’s official “Tor Middle Guard” page [65], requiring thorough site navigation.

The second subtheme is **“Preference for Automation.”** Participants desire automated assistance to ease manual update complexities. While Tor has some automation, operators seek a streamlined setup. A participant stressed this need, stating:

“P30: An easier setup for automated updates on the ‘official’ Tor site. Currently, I manually follow 5 steps, takes 20 minutes to set up...”

4.3.2.3 Tor Auto-Update is Simple

The theme “Tor Auto-Update is Simple” underscores some participants’ belief that the update process is already easy. One participant explicitly stated:

“P7: My Tor relay updates itself whenever there is an update. This is done in the way in which it is described on the Tor website.”

This theme highlights uneven information distribution among relay operators. Considering their varied expertise [67], provided information should bridge this gap for effective support.

4.3.2.4 Update Side Effects

The theme “Update Side Effects” has two subthemes - “Update Penalization Concerns” and “Machine Stability Concerns,” reflecting relay operators’ worries about updates’ side effects on their relays.

The subtheme **“Update Penalization Concerns”** reveals relay operators’ worries about updating impacting their relay stability status. One participant explicitly mentioned this concern:

“P38: ...the network’s routing algorithms seemed to penalise those who updated Tor (and other parts of their systems) regularly...”

The second subtheme **“Machine Stability Concerns”** focuses on relay operators’ desire for updates to not harm their systems’ stability. They prefer *“P41: lightweight”*

Table 4.7: Subthemes informing each theme

Theme	Subthemes	Frequency
Support for New Auto-Updates		3
Desire to Opt-Out of Auto-Updates		5
Recommendations for Implementation	•Accessibility	4
	•Update Timing	7
Opposition to the New Auto-Updates	•Satisfaction with Current Updating	4
	•Concerns Over New Auto-Updates	23

updates to avoid crashes and minimize risks to machine stability. This subtheme also emphasizes concerns about forced updates and associated risks, highlighting the significance of maintaining control over the update process:

“P46: ...ensure tor relay operators are never forced to accept an update... Every node has unique hardware, software configuration, and non-technical conditions...”

4.3.2.5 Update as Instructed by Tor

The theme “Update as Instructed by Tor” reflects relay operators who prioritize updates based on Tor’s recommendations over personal factors. One participant explicitly stated:

“P50: ...Updating is very important and personal considerations (is) low...”

This factor differs from previous research, possibly stemming from relay operators’ acknowledgement of Tor’s sensitivity and the responsibility they hold [32]. Cultivating this compliance mindset could strengthen relay updates.

4.3.3 Analyzing Feedback on Automatic Updates

The thematic analysis of the responses to Question 15 revealed six subthemes, which were then grouped into four key themes: Support for New Auto-Updates, Desire to Opt-Out of Auto-Updates, Recommendations for Implementation, and Opposition to the New Auto-Updates (Table 4.7). Frequency indicates theme prevalence. *These themes reveal participant views and influential factors toward auto-updates, distinguishing them from the existing manual update context in Tor.* This informs Objectives 2 and 3. The complete list of coded responses can be found in Appendix E.2.

4.3.3.1 Support for New Auto-Updates

The “Support for New Auto-Updates” theme includes participants favouring the new automatic update mechanism due to its potential to reduce perceived costs, such as time. One participant mentioned,

“P37: ...Bring it on. One thing less for me to have to be concerned with.”

4.3.3.2 Desire to Opt-Out of Auto-Updates

The “Desire to Opt-Out of Auto-Updates” theme reflects participants who prefer the choice to avoid automatic updates. While not necessarily opposed to the mechanism, they value the control to manually update. Participants explicitly expressed the need for an opt-out choice:

“P10: It should be opt-in only. Or at least opt-out should be possible.”

4.3.3.3 Recommendations for Implementation

The theme “Recommendations for Implementation” consists of two subthemes: “Accessibility” and “Update Timing.” This theme provides implementation insights that could impact the effectiveness of the new automatic update framework.

The “**Accessibility**” subtheme underscores Tor relay diversity across operating systems. This concern aligns with manual update findings, contributing to Objectives 2 and 3. Specific to Objective 2, another contributing factor involves specific considerations, including language barriers and diverse operating practices. A participant reported:

“P48: The automatic update process will need to consider network diversity...”

The concern about the language barrier is valid given the global diversity of Tor relay operators [70]. And as one participant stated,

“P22: Not all Tor relay operators speak English...is there some way to translate their comment into English and vice versa if they want to read protests made in English...”

The second subtheme “**Update Timing**” centres on participants’ concerns about simultaneous node updates causing downtime and network unavailability. This factor emerged from the survey’s automatic update framework. However, it applies broadly to other potential automatic update systems aiming to synchronize relay updates. This demands careful consideration of update timing, as high simultaneous update volume could significantly impact network traffic [28]. One participant specifically points out:

“P16: ...If all servers are upgrading at the same time and there is a problem, a large part of the Tor network will be offline...”

4.3.3.4 Opposition to the New Auto-Updates

The theme “Opposition to New Auto-Updates” includes participants resisting the new automatic update mechanism in Tor relays. This opposition is seen through “Satisfaction with Current Updating” and “Concerns Over New Auto-Updates”.

The subtheme “**Satisfaction with Current Updating**” captures participants’ contentment with the existing Tor relay update process. They express no desire for a new mechanism and share positive experiences confirming its functionality:

“P7: ...tor updates periodically and keeping everything up (to) date to keep the whole thing low maintenance... Never had any troubles with running the system...”

In the second subtheme, “**Concerns Over New Auto-Updates,**” participants voice reservations about the automatic update mechanism. Their varied concerns contribute to Objective 2 by providing specific considerations for updates. These concerns, distinct from previous research and Question 10 analysis, *offer valuable insights into the issues an automatic update mechanism should address.* A recurring view from previous analysis is reiterated here, emphasizing the significance of relay operators’ control over their operational relays. One participant highlighted:

“...As a Tor operator I’m responsible for and in control of my hardware, software and network...”

Furthermore, certain participants believe the burden on relay operators is too high, aligning with the Q10 analysis. This further underscores the need to account for the diverse operator knowledge levels:

“P39: ...as a simple relay operator you probably don’t have the overview of fine-grained implementations of certain goals (e.g. ”tamper protection“)...”

Another concern arises from the analysis, which is not identified in any previous research. It relates to the current low participation rate in the Tor forum, suggesting scepticism about the mechanism’s effectiveness:

“P30: Due to low participation in the community, it’s likely no one will check the update during protest period... ”

Fundamentally, responses stressed that the mechanism could pose challenges to network maintenance. This should be factored into the automatic update mechanism design, ensuring updates remain manageable—a point not covered in prior research:

“P32: ...Do we want unattended, zombie relays to continue transiting traffic if they’re not running on systems which are actively maintained?”

Chapter 5

Discussion

This section outlines practical design recommendations drawn from my study's findings, offering insights for future investigations into improved automatic update mechanisms, addressing Objective 4. Some of these recommendations are applicable beyond Tor, while others are Tor-specific.

5.1 Security Considerations

In my study, security emerged as a prevalent concern for Tor relay operators, consistent with prior research[24, 50, 51]. This emphasizes the need for careful security considerations when designing automatic relay update frameworks. Quantitative analysis showed participants highly prioritized security-related updates and expressed concerns about bugs or malicious code, underscoring security's crucial role in their decision-making. Moreover, participants' attitudes towards automatic updates are more cautious when compared to manual updates. The introduction of automatic updates intensifies their concern for security. This cautious approach aligns with findings from prior research, which also highlighted individuals' reservations about the security implications of automatic updates [5, 22, 87]. The qualitative data analysis also supported it.

Therefore, investing in security is a potential strategy to encourage more relay operators to update their relays. It is vital to address potential attacks on the mechanism, as identified in the qualitative analysis. To achieve this, update mechanisms should emphasize and advertise their security features, making potential operators more aware of the security enhancements. Simultaneously, update mechanisms must avoid adding complexities to relay maintenance.

Nonetheless, a subset of participants conveyed satisfaction with the existing update

design, particularly in relation to the concept of End-of-Life (EOL) for Tor relays. These individuals contend that no additional update mechanism is necessary. Yet, the EOL approach in Tor may not be fully secure, as the protocol's tolerance for multiple versions within the network can leave it vulnerable to attacks [38, 60, 74]. While I did not conduct an in-depth evaluation of participants' perceptions of EOL in Tor relays, future work could focus on designing more efficient EOL processes that consider both the perspectives of relay operators and developers.

5.2 Reduce Perceived Cost

Minimizing perceived costs is another crucial factor in encouraging relay operators to update their relays, aligning with prior studies [51, 55, 71]. While the quantitative analysis showed that updating time was considered of low importance among all 54 participants, 2 out of the 4 participants who did not update mentioned the time-consuming nature of manual updates. Concerns about compatibility issues and potential machine breakdowns due to updates were also identified. The key takeaway from these results is that update designers should ensure that individuals who update their relays do not risk losing access to their operating machines, although guaranteeing 100% accuracy can be challenging [24]. To address this, new relay updates could provide assurance to alleviate their worries.

Furthermore, a unique perceived cost within Tor was identified through qualitative analysis. Unlike previous research findings, participants believe that the Tor network appears to prioritize relays with higher up-time, potentially penalizing those who update frequently. This perception could cause relay operators to postpone updates until major versions are available and maintain longer uptimes. Consequently, when formulating automatic update strategies, it is imperative to address this apprehension.

5.3 Balance Relay Operator Control

Furthermore, achieving an appropriate balance of control is paramount, particularly within the unique context of Tor. The current level of control outlined in the surveyed auto-update mechanism appears insufficient for relay operators' needs. This observation underlines the necessity for thoughtful technical development. Similarly, participant feedback indicates that it's crucial not to overload relay operators with excessive technical responsibilities. These operators are volunteers with varying backgrounds

and expertise in relay operations [37]. Imposing excessive duties could lead to pressure and potentially result in updates being pushed without a genuine community consensus. It's imperative to ensure that both the community and updates are guided by active and diverse participation, rather than a select few.

5.4 Improve Update Communication

Effective communication is also crucial when designing update mechanisms, and several studies have emphasized the significance of clear and transparent update messages [23, 26, 33, 34, 35, 82]. Although annoyance was not a prevalent factor in this survey, confusion was clearly observed both before and after updates.

In the quantitative analysis, the factor "I do not keep up with Tor update news regularly" was rated as of low importance, which is a factor related to communication. However, in the qualitative part, participants provided several new communication-related factors that are specifically applicable to the Tor context. One crucial communication aspect desired by participants is an update road map, suggesting what fixes are on the way. This approach has several advantages and is widely accepted in the product management world [54]. It could provide transparency [10], manage expectations for what developers are building, and create more opportunities for user feedback [6]. Despite the challenges that this approach may present [54], it is crucial to include it in the relay update mechanism. Conducting future research on how users feel about this approach and whether it can offer a clearer and more comprehensive picture to relay operators is of utmost importance.

Moreover, the qualitative analysis indicated a need for more easily accessible and clear notifications. While previous research has also highlighted this necessity [41], what sets this study apart is the fact that not all relay operators subscribe to relay operators' lists for prompt notifications. To ensure relay operators are well-informed about available updates, clear notifications through alternative channels should be explored. Similarly, post-update notifications were also requested during the survey to enable relay operators to receive prompt update results. Future research can explore effective notification methods for delivering clear and timely information before and after updates. Such communication improvements can enhance operator confidence and update efficiency.

5.5 Platform Inconsistency Need to be Addressed

Platform inconsistency within the Tor network is another factor that requires specific attention. This is a topic not explored in previous research and it specifically fits in the context of Tor. The qualitative data collected revealed a theme related to platform inconsistency, indicating the concern relay operators have about platforms that are behind in updates, especially when security fixes are urgently needed. Given that Tor relays can operate on diverse platforms [65], update synchronization discrepancies can arise. To ensure an effective update process, it is essential to design an automatic update mechanism that can accommodate various platforms and efficiently deliver updates to all relays when required. By doing so, the overall security of the network will be improved, and the risk of relays falling behind due to platform inconsistency will be minimized.

5.6 Limitations

The present study comes with certain limitations. One noteworthy limitation pertains to the distribution of participants. The demographics of the participant pool may not offer a fully representative view of the entire relay operators. This situation arises from the existing gap in research within this specific domain, which positions this study as an exploratory step into investigating user attitudes within the Tor environment. Moreover, as inferred from the analysis, the viewpoints of the present participants already display a certain degree of diversity. Consequently, a broader and more comprehensive participant pool would be prudent to enhance the broader applicability of the conclusions drawn. Notably, the current sample size, although not representative, is comprised of individuals who display a strong concern for security within the Tor community. As seen in the Literature Review, those actively involved in forums and feedback provision often influence a broader audience with their perspectives. Consequently, their attitudes can influence a wider audience, making their opinions highly valuable.

Furthermore, the limitations associated with the survey method itself need consideration. The specific nature of Tor's operation, coupled with the survey's design centred around the automatic update mechanism, implies that participants did not have the opportunity to tangibly interact with the actual implementation. To address this, a follow-up study could be conducted. By incorporating participants' suggestions into the update design and subsequently conducting interviews, we could gauge how their

experience with the new design compares to the original one. Ideally, a field study could be undertaken, allowing participants to actively engage with the design.

Another limitation is in the questionnaire design regarding “BandwidthRate.” Expert consultations prior to the survey revealed concerns that seeking specific bandwidth usage data could risk deanonymization and deter participants. To address this, the survey focused on relay operators’ perception of capacity (BandwidthRate), not precise usage statistics. However, in reviewing the responses, it became evident that some participants had provided actual bandwidth usage data instead of the configured BandwidthRate value. Due to the prevalence of this misunderstanding, the results of this specific question cannot be utilized in the analysis. Upon further consultation with an experienced relay operator, it seems that many operators do not even set this value, relying instead on the “Advertised Bandwidth” measurement [68]. The misunderstanding highlights the need for clearer guidance and support materials. Future research should explore more reliable methods for assessing bandwidth capabilities, like Advertised Bandwidth. Given some participants’ willingness to share actual usage statistics, surveying real bandwidth data could definitely provide useful insights, with appropriate considerations for privacy. More broadly, the issues around BandwidthRate signify the importance of establishing clear definitions and shared understandings of technical parameters when conducting research with the Tor community. This presents an opportunity to reevaluate which metrics are most meaningful and improve educational resources accordingly.

Lastly, it is important to acknowledge a limitation that is common in security research, namely the reliance on self-reported behaviours. This limitation is shared with various previous studies [83]. However, in recognition of this potential limitation, this study sought to improve the robustness of its analysis. This was accomplished by employing a comprehensive approach that include both quantitative and qualitative data.

Chapter 6

Conclusion

This research aimed to advance understanding of user attitudes toward automatic updates within the Tor environment. The specific objectives were to:

1. *Identify* established factors impacting user attitudes towards updates.
2. *Explore* Tor relay operators' views and practices related to updating relays.
3. *Examine* how relay operators' attitudes differ or change in the context of manual and automatic updates.
4. *Formulate recommendations* on how automatic updates should be designed.

This concluding chapter will revisit these research objectives, summarize the study's key findings, and offer conclusions based on the results. Future research directions will be proposed to progress this work further. By adopting this structure, the research is concluded by reflecting on whether the initial aims were achieved, including assessing the overall value of the study.

6.1 Conclusion

6.1.1 Objective 1: Update Drivers and Barriers

The literature identified potential reasons for varied user attitudes and behaviours toward updating. On the drivers' side, it became evident that relay operators' interpretation of update messages is highly influential in shaping their actions. The study highlighted the impact of crowd-sourcing, where users frequently seek advice from their peers when confronted with update-related challenges. Additionally, the perceived necessity of

an update emerged as a driving factor. Security considerations also played a pivotal role in influencing users' attitudes towards updates. Moving to the barriers, a central finding was the challenge users face in fully comprehending the messages conveyed with updates. An important factor contributing to this complexity is the emotion of annoyance, with intrusive update messages deterring users from promptly engaging with the updates. Moreover, users' specific characteristics and technical expertise play significant roles in shaping their attitudes toward updates. Notably, this exploration uncovered the variability of user attitudes across different contexts, underscoring the significance of contextual specificity in understanding user behaviours and decision-making processes related to updates. The findings from this exploration shed light on the intricate dynamics that shape relay operators' update decisions, contributing to a deeper understanding of user attitudes within the realm of update mechanisms.

6.1.2 Objective 2: Tor relay Views and Practices

This empirical research identified several factors previously found to influence user updating that proved inapplicable in the Tor context. For example, necessity or importance perceptions were less salient for operators. The findings suggest relay operators take a more objective view of updates, prioritizing technical implications like security value over subjective judgments. In addition, new influential factors also emerged, including preferences for control and frustrations with platform inconsistency, identifying context-specific dynamics. Critically, with a diverse sample of only 54 participants, no singular perspective dominated - operator views varied substantively. In summary, this exploration uncovers shared and distinct updating attitudes between the general context and the Tor operator community. While some common beliefs have limited relevance, new motivations arise from this distinct technical environments. Further research is essential to capture the full breadth of this complex perspective spectrum.

6.1.3 Objective 3: Manual Update versus Automatic Update

This research conducted a direct comparison between operator attitudes concerning manual and automatic update methods. The findings revealed that when participants were presented with the option of an automated approach, they exhibited greater overall apprehension, even though they acknowledged potential benefits such as more efficient workflows. Notably, concerns centred around critical risks associated with automation, such as the potential for single points of failure and a decrease in operator control.

While general attitudes remained relatively consistent across different factors, negative responses were notably amplified for automation drawbacks. This divergence underscores the influence of update context on perspectives. Operators exhibited a higher level of cautiousness toward automated methods compared to manual ones. Overall, this variability in attitudes highlights the necessity of considering both advantages and risks when designing automatic update systems for this community.

6.1.4 Objective 4: Recommendations for Automatic Update Design

Having fulfilled Objectives 1 to 3, Objective 4 was subsequently accomplished by providing recommendations in the discussion section of the report. These recommendations focus on factors that should be carefully considered when designing automatic updates for Tor relays. These factors encompass the importance of security considerations, the reduction of perceived costs for relay operators, the need to balance control between operators and automation, enhancements in communication strategies for updates, and addressing the issue of platform consistency in Tor relays. These suggestions are novel and systematically structured, reflecting the considerations of relay operators in the design of automatic update for the Tor.

6.2 Future work

While this research has successfully achieved its goal of gaining deeper insights into relay operators' attitudes towards automatic updates, certain limitations should be acknowledged. Firstly, the survey conducted is not fully representative of the entire population of Tor relay operators. This limitation restricts the generalizability of the findings to the whole relay community. As further research delves into these issues and conducts more comprehensive case studies, an enhanced understanding of the topics is expected to develop, contributing to the broader landscape of relay operators' attitudes research.

Another potential research direction is utilizing more practical methods beyond descriptive questionnaires to address self-reporting limitations. For example, conducting field studies would allow participants to directly interact with an update framework. By using this approach, researchers could observe how people behave and listen to their opinions in a real-world setting. This could provide a better understanding of what people do, why they do it, and what obstacles they encounter, all in a more natural

and practical environment. Interviews during field testing could also capture detailed perspectives. Hands-on engagement provides a fuller sense of factors influencing behaviours. While questionnaires offer breadth, complementing with immersive studies enhances depth. A multifaceted methodology combining surveys, field research, and interviews may provide the most accurate representation of operator patterns around updating.

A third direction for future exploration lies in the temporal aspect. This empirical study serves as an initial attempt to understand relay operators' attitudes towards updates. Conducting a study that spans different time frames could provide a more comprehensive and representative view, showcasing the evolution of activities and decision-making processes. However, due to the constraints of the dissertation's timeline, such an approach was not feasible. Implementing such longitudinal studies could yield valuable insights in the future.

Bibliography

- [1] Microsoft 365. The importance of computer software updates security patches, Jul 2022. URL <https://www.microsoft.com/en-us/microsoft-365-life-hacks/privacy-and-safety/computer-software-update>. Accessed 16th August 2023.
- [2] Agrawal Abhishek, Fantham David, Debraj Ghosh, et al. Microsoft security intelligence report volume 24. *Microsoft security intelligence report*, pages 20–24, 2019. URL <https://clouddamcdnprodep.azureedge.net/gdc/gdc09FrGq/original>. Accessed 16th August 2023.
- [3] I. Elaine Allen and Christopher A. Seaman. Likert scales and data analyses. *Quality progress*, 40(7):64–65, 2007.
- [4] Ron Amadeo. Adware vendors buy chrome extensions to send ad- and malware-filled updates, Jan 2014. URL <https://arstechnica.com/information-technology/2014/01/malware-vendors-buy-chrome-extensions-to-send-adware-filled-updates/>. Accessed 16th August 2023.
- [5] Lisanne Bainbridge. Ironies of automation. In *Analysis, design and evaluation of man-machine systems*, pages 129–135. Elsevier, 1983.
- [6] Jake Bartlett. Should you share your product roadmap publicly?, Nov 2020. URL <https://www.launchnotes.com/blog/should-you-share-your-product-roadmap-publicly#5-advantages-of-publicly-sharing-your-roadmap>. Accessed 16th August 2023.
- [7] Lamiaa Basyoni, Noora Fetais, Aiman Erbad, Amr Mohamed, and Mohsen Guizani. Traffic analysis attacks on tor: a survey. In *2020 IEEE International Conference on Informatics, IoT, and Enabling Technologies (ICIoT)*, pages 183–188. IEEE, 2020.

- [8] Kevin Bauer, Damon McCoy, Dirk Grunwald, Tadayoshi Kohno, and Douglas Sicker. Low-resource routing attacks against tor. In *Proceedings of the 2007 ACM workshop on Privacy in electronic society*, pages 11–20, 2007.
- [9] Leyla Bilge and Tudor Dumitraş. Before we knew it: an empirical study of zero-day attacks in the real world. In *Proceedings of the 2012 ACM conference on Computer and communications security*, pages 833–844, 2012.
- [10] The Trymata blog. Product roadmap: How it affects user experience, Jun 2023. URL <https://trymata.com/blog/learn/product-roadmap-and-user-experience/>. Accessed 16th August 2023.
- [11] Harry N. Boone Jr and Deborah A. Boone. Analyzing likert data. *The Journal of extension*, 50(2):48, 2012.
- [12] Karoline Busse, Julia Schäfer, and Matthew Smith. Replication: No one can hack my mind revisiting a study on expert and non-expert security practices and advice. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*, pages 117–136, 2019.
- [13] David Byrne. A worked example of braun and clarke’s approach to reflexive thematic analysis. *Quality & quantity*, 56(3):1391–1412, 2022.
- [14] National Cyber Security Centre. Keeping devices and software up to date, June 2021. URL <https://www.ncsc.gov.uk/collection/device-security-guidance/managing-deployed-devices/keeping-devices-and-software-up-to-date>. Accessed 16th August 2023.
- [15] Joseph Check and Russell K. Schutt. *Research methods in education*. Sage publications, 2011.
- [16] Erika Chin, Adrienne Porter Felt, Vyas Sekar, and David Wagner. Measuring user confidence in smartphone security and privacy. In *Proceedings of the eighth symposium on usable privacy and security*, pages 1–16, 2012.
- [17] Ben Derrick and Paul White. Comparing two samples from an individual likert question. *International Journal of Mathematics and Statistics*, 18(3):1–13, 2017.
- [18] Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. In *Proceedings of the 13th Conference on USENIX*

- Security Symposium - Volume 13*, SSYM'04, page 21, USA, 2004. USENIX Association.
- [19] Thomas Duebendorfer and Stefan Frei. Why silent updates boost security. *TIK, ETH Zurich, Tech. Rep*, 302:98, 2009.
- [20] Arthur Edelstein. Publish tor relay auto-update instructions, Dec 2017. URL <https://gitlab.torproject.org/tpo/web/trac/-/issues/24505>. Accessed 16th August 2023.
- [21] Thomas W. Edgar and David O. Manz. Chapter 4 - exploratory study. In *Research Methods for Cyber Security*, pages 95–130. Syngress, 2017.
- [22] W. Keith Edwards, Erika Shehan Poole, and Jennifer Stoll. Security automation considered harmful? In *Proceedings of the 2007 Workshop on New Security Paradigms, White Mountain Hotel and Resort, New Hampshire, USA - September 18-21, 2007*, pages 33–42. ACM, 2007.
- [23] Michael Fagan, Mohammad Maifi Hasan Khan, and Ross Buck. A study of users' experiences and beliefs about software update messages. *Computers in Human Behavior*, 51:504–519, 2015.
- [24] Sadeqh Farhang, Jake Weidman, Mohammad Mahdi Kamani, Jens Grossklags, and Peng Liu. Take it or leave it: A survey study on operating system upgrade practices. In *Proceedings of the 34th Annual Computer Security Applications Conference, ACSAC '18*, page 490–504. Association for Computing Machinery, 2018.
- [25] Adrienne Porter Felt, Erika Chin, Steve Hanna, Dawn Song, and David Wagner. Android permissions demystified. In *Proceedings of the 18th ACM conference on Computer and communications security*, pages 627–638, 2011.
- [26] Adrienne Porter Felt, Elizabeth Ha, Serge Egelman, Ariel Haney, Erika Chin, and David Wagner. Android permissions: User attention, comprehension, and behavior. In *Proceedings of the eighth symposium on usable privacy and security*, pages 1–14, 2012.
- [27] Eric Filiol, Maxence Delong, and J. Nicolas. Statistical and combinatorial analysis of the TOR routing protocol: structural weaknesses identified in the TOR network. *J. Comput. Virol. Hacking Tech.*, 16(1):3–18, 2020.

- [28] Electronic Frontier Foundation, Feb 2023. URL <https://www.eff.org/pages/what-tor-relay#:~:text=The%20Tor%20software%20depends%20on,people%20to%20run%20Tor%20relays>. Accessed 16th August 2023.
- [29] Alexander Færøy. Tor relay guide, Jun 2020. URL <https://gitlab.torproject.org/legacy/trac/-/wikis/TorRelayGuide#AutomaticSoftwareUpdates>. Accessed 16th August 2023.
- [30] GeKo. [tor-relays] psa: Tor 0.4.5 reaches end of life (eol) on 2023-02-15, Jan 2023. URL <https://forum.torproject.org/t/tor-relays-psa-tor-0-4-5-reaches-end-of-life-eol-on-2023-02-15/6338>. Accessed 16th August 2023.
- [31] Jeremy Goecks and Elizabeth D Mynatt. Social approaches to end-user privacy management. *Security and Usability: Designing Secure Systems That People Can Use*, pages 523–545, 2005.
- [32] Gus. Expectations for relay operators, Apr 2020. URL <https://gitlab.torproject.org/tpo/community/team/-/wikis/Expectations-for-Relay-Operators>. Accessed 16th August 2023.
- [33] Marian Harbach, Sascha Fahl, Thomas Muders, and Matthew Smith. Towards measuring warning readability. In *Proceedings of the 2012 ACM conference on Computer and communications security*, pages 989–991, 2012.
- [34] Marian Harbach, Sascha Fahl, Polina Yakovleva, and Matthew Smith. Sorry, I don't get it: An analysis of warning message texts. In *Financial Cryptography and Data Security - FC 2013 Workshops, USEC and WAHC 2013, Okinawa, Japan, April 1, 2013, Revised Selected Papers*, volume 7862 of *Lecture Notes in Computer Science*, pages 94–111. Springer, 2013.
- [35] Marian Harbach, Markus Hettig, Susanne Weber, and Matthew Smith. Using personal examples to improve risk communication for security & privacy decisions. In *CHI Conference on Human Factors in Computing Systems, CHI'14, Toronto, ON, Canada - April 26 - May 01, 2014*, pages 2647–2656. ACM, 2014.
- [36] Ayako Akiyama Hasegawa, Naomi Yamashita, Tatsuya Mori, Daisuke Inoue, and Mitsuaki Akiyama. Understanding non-experts' security- and privacy-related questions on a q&a site. In *Eighteenth Symposium on Usable Privacy and Security, SOUPS 2022, Boston, MA, USA, August 7-9, 2022*, pages 39–56. USENIX Association, 2022.

- [37] Hsiao-Ying Huang and Masooda Bashir. Who is behind the onion? understanding tor-relay operators. URL https://cups.cs.cmu.edu/soups/2015/posters/soups2015_posters-final12.pdf. Accessed 16th August 2023.
- [38] Rob Jansen, Tavish Vaidya, and Micah Sherr. Point break: A study of bandwidth denial-of-service attacks against tor. In *28th USENIX Security Symposium, USENIX Security 2019, Santa Clara, CA, USA, August 14-16, 2019*, pages 1823–1840. USENIX Association, 2019.
- [39] Aaron Johnson, Chris Wacek, Rob Jansen, Micah Sherr, and Paul F. Syverson. Users get routed: traffic correlation on tor by realistic adversaries. In *2013 ACM SIGSAC Conference on Computer and Communications Security, CCS'13, Berlin, Germany, November 4-8, 2013*, pages 337–348. ACM, 2013.
- [40] Ishan Karunanayake, Nadeem Ahmed, Robert A. Malaney, Rafiqul Islam, and Sanjay K. Jha. De-anonymisation attacks on tor: A survey. *IEEE Commun. Surv. Tutorials*, 23(4):2324–2350, 2021.
- [41] Patrick Gage Kelley, Lorrie Faith Cranor, and Norman M. Sadeh. Privacy as part of the app decision-making process. In *2013 ACM SIGCHI Conference on Human Factors in Computing Systems, CHI '13, Paris, France, April 27 - May 2, 2013*, pages 3393–3402. ACM, 2013.
- [42] Moazzam Khan, Zehui Bi, and John A. Copeland. Software updates as a security metric: Passive identification of update trends and effect on machine infection. In *31st IEEE Military Communications Conference, MILCOM 2012, Orlando, FL, USA, October 29 - November 1, 2012*, pages 1–6. IEEE, 2012.
- [43] John Knight. Tor browser bundle-tor goes portable, Sep 2011. URL <https://www.linuxjournal.com/content/tor-browser-bundle-tor-goes-portable>. Accessed 16th August 2023.
- [44] Oscar Labra, Carol Castro, Robin Wright, and Isis Chamblás. *Thematic Analysis in Social Work: A Case Study*, pages 183–202. 12 2019.
- [45] Rensis Likert. A technique for the measurement of attitudes. *Archives of psychology*, 1932.

- [46] Alexander De Luca, Sauvik Das, Martin Ortlieb, Iulia Ion, and Ben Laurie. Expert and non-expert attitudes towards (secure) instant messaging. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, pages 147–157, 2016.
- [47] Chris Manning. Maintenance matters: Timely upgrades, Mar 2023. URL <https://www.viget.com/articles/maintenance-matters-timely-upgrades/>. Accessed 16th August 2023.
- [48] Géraldine Vache Marconato, Vincent Nicomette, and Mohamed Kaâniche. Security-related vulnerability life cycle analysis. In *7th International Conference on Risks and Security of Internet and Systems, CRiSIS 2012, Cork, Ireland, October 10-12, 2012*, pages 1–8. IEEE Computer Society, 2012.
- [49] Arunesh Mathur and Marshini Chetty. Impact of user characteristics on attitudes towards automatic mobile application updates. In *Thirteenth Symposium on Usable Privacy and Security, SOUPS 2017, Santa Clara, CA, USA, July 12-14, 2017*, pages 175–193. USENIX Association, 2017.
- [50] Arunesh Mathur, Josefine Engel, Sonam Sobti, Victoria Chang, and Marshini Chetty. “they keep coming back like zombies”: Improving software updating interfaces. In *Twelfth Symposium on Usable Privacy and Security, SOUPS 2016, Denver, CO, USA, June 22-24, 2016*, pages 43–58. USENIX Association, 2016.
- [51] Arunesh Mathur, Nathan Malkin, Marian Harbach, Eyal Péer, and Serge Egelman. Quantifying users’ beliefs about software updates. *CoRR*, abs/1805.04594, 2018.
- [52] Gary E. Meek, Ceyhun Ozgur, and Kenneth Dunning. Comparison of the t vs. wilcoxon signed-rank test for likert scale data and small samples. *Journal of modern applied statistical methods*, 6(1):10, 2007.
- [53] Andreas Möller, Florian Michahelles, Stefan Diewald, Luis Roalter, and Matthias Kranz. Update behavior in app markets and security implications: A case study in google play. In *Research in the Large, LARGE 3.0: 21/09/2012-21/09/2012*, pages 3–6, 2012.
- [54] Bri Morgaine. Should you have a public product roadmap?, Jul 2023. URL <https://canny.io/blog/should-you-have-a-public-roadmap/>. Accessed 16th August 2023.

- [55] Jason Morris, Ingolf Becker, and Simon Parkin. An analysis of perceptions and support for windows 10 home edition update features. *J. Cybersecur.*, 6(1), 2020.
- [56] Steven J. Murdoch and George Danezis. Low-cost traffic analysis of tor. In *2005 IEEE Symposium on Security and Privacy (S&P 2005), 8-11 May 2005, Oakland, CA, USA*, pages 183–195. IEEE Computer Society, 2005.
- [57] Milad Nasr, Alireza Bahramali, and Amir Houmansadr. Deepcorr: Strong flow correlation attacks on tor using deep learning. *CoRR*, abs/1808.07285, 2018.
- [58] Kartik Nayak, Daniel Marino, Petros Efstathopoulos, and Tudor Dumitras. Some vulnerabilities are different than others - studying vulnerabilities and attack surfaces in the wild. In *Research in Attacks, Intrusions and Defenses - 17th International Symposium, RAID 2014, Gothenburg, Sweden, September 17-19, 2014. Proceedings*, volume 8688 of *Lecture Notes in Computer Science*, pages 426–446. Springer, 2014.
- [59] Le T. Nguyen, Yuan Tian, Sungho Cho, Wookjong Kwak, S. Parab, Yu Seung Kim, Patrick Tague, and Joy Zhang. Unlocin: Unauthorized location inference on smartphones without being caught. In *2013 International Conference on Privacy and Security in Mobile Systems, PRISMS 2013, Atlantic City, NJ, USA, June 24-27, 2013*, pages 1–8. IEEE, 2013.
- [60] Lasse Øverlier and Paul F. Syverson. Locating hidden servers. In *2006 IEEE Symposium on Security and Privacy (S&P 2006), 21-24 May 2006, Berkeley, California, USA*, pages 100–114. IEEE Computer Society, 2006.
- [61] Pierluigi Paganini. Tor project is going to remove end-of-life relays from the network, Oct 2019. URL <https://securityaffairs.com/92329/security/tor-project-eol-relays.html>. Accessed 16th August 2023.
- [62] Nandita Pattnaik, Shujun Li, and Jason R. C. Nurse. Perspectives of non-expert users on cyber security and privacy: An analysis of online discussions on twitter. *Comput. Secur.*, 125:103008, 2023.
- [63] Jonald L Pimentel. A note on the usage of likert scaling for research data analysis. *USM R&D Journal*, 18(2):109–112, 2010.
- [64] Jon Postel. Transmission control protocol. *RFC*, 793:1–91, 1981.

- [65] The Tor Project. Relay operations - middle/guard relay, . URL <https://community.torproject.org/relay/setup/guard/>. Accessed 16th August 2023.
- [66] The Tor Project. Automatic updates - debian and ubuntu, . URL <https://community.torproject.org/relay/setup/guard/debian-ubuntu/updates/>. Accessed 16th August 2023.
- [67] The Tor Project. Community - relay operations, . URL <https://community.torproject.org/relay/>. Accessed 16th August 2023.
- [68] The Tor Project. Tor metrics glossary, 2023. URL <https://metrics.torproject.org/glossary.html#advertised-bandwidth>. Accessed 16th August 2023.
- [69] Emilee J. Rader, Rick Wash, and Brandon Brooks. Stories as informal lessons about security. In *Symposium On Usable Privacy and Security, SOUPS '12, Washington, DC, USA - July 11 - 13, 2012*, page 6. ACM, 2012.
- [70] Chaitanya Rahalkar, Anushka Virgaonkar, and Kethaki Varadan. Analyzing trends in tor. *CoRR*, abs/2208.11149, 2022.
- [71] Prashanth Rajivan, Efrat Aharonov-Majar, and Cleotilde Gonzalez. Update now or later? effects of experience, cost, and risk preference on update decisions. *J. Cybersecur.*, 6(1):tyaa002, 2020.
- [72] Denise Rey and Markus Neuhäuser. Wilcoxon-signed-rank test. In *International Encyclopedia of Statistical Science*, pages 1658–1659. Springer, 2011.
- [73] Florentin Rochet and Tariq Elahi. Towards flexible anonymous networks. *CoRR*, abs/2203.03764, 2022.
- [74] Florentin Rochet and Olivier Pereira. Dropping on the edge: Flexibility and traffic confirmation in onion routing protocols. *Proc. Priv. Enhancing Technol.*, 2018(2): 27–46, 2018.
- [75] Mehran Sahami, Susan Dumais, David Heckerman, and Eric Horvitz. A bayesian approach to filtering junk e-mail. In *Learning for Text Categorization: Papers from the 1998 workshop*, volume 62, pages 98–105. Citeseer, 1998.
- [76] Niklas Schmucker. Web tracking. In *SNET2 Seminar Paper-Summer Term*, volume 2011. Citeseer, 2011.

- [77] Samuel Sanford Shapiro and Martin B. Wilk. An analysis of variance test for normality (complete samples). *Biometrika*, 52(3/4):591–611, 1965.
- [78] Gail M. Sullivan and Anthony R. Artino Jr. Analyzing and interpreting data from likert-type scales. *Journal of graduate medical education*, 5(4):541–542, 2013.
- [79] CryptPad Team. Cryptpad documentation, 2023. URL <https://docs.cryptpad.org/en/>. Accessed 16th August 2023.
- [80] L. Tenenboim-Chekina, O. Barad, A. Shabtai, D. Mimran, L. Rokach, B. Shapira, and Y. Elovici. Detecting application update attack on mobile devices through network featur. In *2013 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pages 91–92, 2013.
- [81] The Tor Project. Relay operations - types of relays on the tor network. URL <https://community.torproject.org/relay/types-of-relays/>. Accessed 16th August 2023.
- [82] Yuan Tian, Bin Liu, Weisi Dai, Blase Ur, Patrick Tague, and Lorrie Faith Cranor. Supporting privacy-conscious app update decisions with user reviews. In *Proceedings of the 5th Annual ACM CCS Workshop on Security and Privacy in Smartphones and Mobile Devices, SPSM 2015, Denver, Colorado, USA, October 12, 2015*, pages 51–61. ACM, 2015.
- [83] Kami Vaniea and Yasmeen Rashidi. Tales of software updates: The process of updating software. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, San Jose, CA, USA, May 7-12, 2016*, pages 3215–3226. ACM, 2016.
- [84] Kami Vaniea, Emilee J. Rader, and Rick Wash. Betrayed by updates: how negative experiences affect future security. In *CHI Conference on Human Factors in Computing Systems, CHI'14, Toronto, ON, Canada - April 26 - May 01, 2014*, pages 2671–2674. ACM, 2014.
- [85] Nadav Voloch and Maor Meir Hajaj. Handling exit node vulnerability in onion routing with a zero-knowledge proof. In *Information Integration and Web Intelligence - 24th International Conference, iiWAS 2022, Virtual Event, November 28-30, 2022, Proceedings*, volume 13635 of *Lecture Notes in Computer Science*, pages 399–405. Springer, 2022.

- [86] Alex Vovk. The importance of timely patching, Aug 2022. URL <https://securityboulevard.com/2022/08/the-importance-of-timely-patching/>. Accessed 16th August 2023.
- [87] Rick Wash, Emilee J. Rader, Kami Vaniea, and Michelle Rizer. Out of the loop: How automated software updates cause unintended security consequences. In *Tenth Symposium on Usable Privacy and Security, SOUPS 2014, Menlo Park, CA, USA, July 9-11, 2014*, pages 89–104. USENIX Association, 2014.
- [88] Xiong Zhen-hai and Yang Yong-zhi. Automatic updating method based on maven. In *2014 9th International Conference on Computer Science & Education*, pages 1074–1077. IEEE, 2014.
- [89] Yajin Zhou and Xuxian Jiang. Dissecting android malware: Characterization and evolution. In *IEEE Symposium on Security and Privacy, SP 2012, 21-23 May 2012, San Francisco, California, USA*, pages 95–109. IEEE Computer Society, 2012.

Appendix A

Combined Participants' Information Sheet and Consent Form

Participant Information Sheet

Project title:	User study of Tor relay operators attitudes towards automatic update
Principal investigator:	Tariq Elahi
Researcher collecting data:	For this project, the researcher will be collecting data through an online survey. The specific data collected will include participants' responses to close and open-ended questions about their beliefs regarding automatic updates in the Tor network. Participants will be asked to indicate how often they have experienced each statement to be true, providing insights into their perspectives and experiences.

This study was certified according to the Informatics Research Ethics Process, reference number 775342. Please take time to read the following information carefully. You should keep this page for your records.

Who are the researchers?

Meitong Wang, Tariq Elahi

What is the purpose of the study?

The goal of this research is to investigate the feasibility and acceptance of an automatic update mechanism for the Tor network. The study will involve conducting surveys with Tor relay operators to gather their opinions on the proposed mechanism. The collected data will be analysed to gain insights into attitudes towards automatic updates, encompassing both favourable and unfavourable perspectives. This comprehensive understanding will enable us to identify potential concerns and areas for improvement in the update process. The results of this research can be used to enhance the security and privacy of the Tor network.



Figure A.1: Page 1 of Participant Information Sheet

Why have I been asked to take part?

You have been asked to take part in this research study because you meet the criteria for being a Tor relay operator. Your participation in the study is valuable because it aims to investigate the attitudes and perspectives of Tor relay operators towards automatic updates. By sharing your experiences and opinions, you can contribute to the understanding of the challenges and concerns related to automatic updates in the Tor network. Your input will help inform the design and improvement of update processes that better address the needs and preferences of relay operators, ultimately enhancing the security and privacy of the Tor network. Your voluntary participation is greatly appreciated and will make a meaningful contribution to the research.

Do I have to take part?

No – participation in this study is entirely up to you. You can withdraw from the study at any time, without giving a reason. Your rights will not be affected. If you wish to withdraw, contact the PI. We will stop using your data in any publications or presentations submitted after you have withdrawn consent. However, we will keep copies of your original consent, and of your withdrawal request.

What will happen if I decide to take part?

The data being collected will include your responses to a survey questionnaire specifically designed for this study. The questionnaire will include questions related to your attitudes, beliefs, and experiences regarding automatic updates in the Tor network. It may cover topics such as concerns or benefits associated with updates, and your overall perspective on automatic update processes.

The data will be collected through an online survey. You will be provided with a link to the survey, and you can access and complete it at your convenience.

The duration of your participation will depend on your own pace and the time it takes for you to complete the survey. It is anticipated that the survey will take approximately 10-20 minutes to finish.

In this particular study, no audio or video recording of participants will take place. The data collection solely relies on the completion of the online survey questionnaire.



Figure A.2: Page 2 of Participant Information Sheet

You will be requested to complete the survey once. You can access the survey using the provided link at any suitable location and time of your choosing. The flexibility allows you to participate in the study at your convenience.

Are there any risks associated with taking part?

There are no significant risks associated with participation.

Are there any benefits associated with taking part?

No.

What will happen to the results of this study?

The results of this study may be summarised in published articles and reports. Quotes or key findings will be anonymized: We will remove any information that could, in our assessment, allow anyone to identify you. With your consent, information can also be used for future research. Your data may be archived for a maximum of two years.

Data protection and confidentiality.

Your data will be processed in accordance with Data Protection Law. All information collected about you will be kept strictly confidential. Your data will be referred to by a unique participant number rather than by name. Your data will only be viewed by the researcher/research team Meitong Wang and Tariq Elahi.

All electronic data will be stored on a password-protected encrypted computer, on the School of Informatics' secure file servers, or on the University's secure encrypted cloud storage services (DataShare). Your consent information will be kept separately from your responses in order to minimise risk.

What are my data protection rights?

The University of Edinburgh is a Data Controller for the information you provide. You have the right to access information held about you. Your right of access can be exercised in accordance Data Protection Law. You also have other rights including rights of correction, erasure and objection. For more details, including the right to lodge a complaint with the Information Commissioner's Office, please visit



Figure A.3: Page 3 of Participant Information Sheet

www.ico.org.uk. Questions, comments and requests about your personal data can also be sent to the University Data Protection Officer at dpo@ed.ac.uk. For general information about how we use your data, go to: edin.ac/privacy-research

Who can I contact?

If you have any further questions about the study, please contact the lead researcher, Meitong Wang at s2447273@ed.ac.uk.

If you wish to make a complaint about the study, please contact inf-ethics@inf.ed.ac.uk. When you contact us, please provide the study title and detail the nature of your complaint.

Updated information.

If the research project changes in any way, an updated Participant Information Sheet will be made available on <http://web.inf.ed.ac.uk/infweb/research/study-updates>.

Consent

By proceeding with the study, I agree to all of the following statements:

- I have read and understood the above information.
- I understand that my participation is voluntary, and I can withdraw at any time.
- I consent to my anonymised data being used in academic publications and presentations.
- I allow my data to be used in future ethically approved research.

Once you have reviewed the Participant Information Sheet and are ready to proceed with the survey, please click on the following link: [\[Take me to the survey\]](#).



Figure A.4: Page 4 of Participant Information Sheet

Appendix B


Questionnaire


This form will close on 7/23/2023, 11:55:00 PM


Tor Relay Update Attitudes and Factors Survey


Responses to this form are anonymized


 Join Our Survey on Automatic Updates for the Tor Network! 

Are you a Tor relay operator? We need your valuable input! Participate in our research study to explore the feasibility and acceptance of an automatic update mechanism for the Tor network. Your opinion matters in shaping the future of Tor updates. 

By completing a simple survey, you can contribute to enhancing the security and privacy of the Tor network. We aim to understand your perspectives on automatic updates, and find ways to improve the update process. Your feedback will help us make the Tor network more secure and reliable for all users. 

Who can benefit from this research? The Tor network itself, the relay operator community, academic researchers focusing on privacy and security, and individuals or organizations relying on the Tor network for online privacy. Your participation will directly contribute to the betterment of this vital network. 

Rest assured, your privacy is of the utmost importance to us. We will only collect your survey responses, without any personally identifiable information. Your anonymity is guaranteed, and your insights will be treated with the utmost confidentiality. 

Join us now in shaping the future of the Tor network's automatic updates! Your opinion matters. Together, let's make Tor even better. 


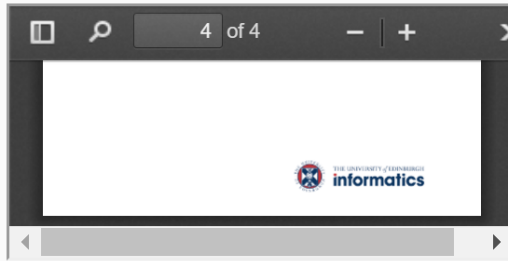
Thank you for your valuable contribution! 

Figure B.1: Section 1 - Survey introduction

Before you begin the survey, we kindly request that you take a moment to read the Participant Information Sheet (PIS). The PIS provides important information about the purpose of this research, the procedures involved, and your rights as a participant. It will help you understand the nature of the study and make an informed decision about your participation.

Once you open the PDF file, you can view it in "presentation mode" or download it using the options available in the top right corner of the file called "Tools."

Once you have familiarized yourself with the information provided in the Participant Information Sheet, you may proceed with the survey. Thank you for your attention and cooperation.



Demographic Questions

1. For how long have you been utilizing Tor as a means to safeguard your online privacy and anonymity?

Required

- Less than 1 year
- 1-3 years
- 3-5 years
- More than 5 years

Figure B.2: Section 1

2. How long have you been a Tor relay operator? Required

Less than 1 year
 1-3 years
 3-5 years
 More than 5 years

3. How many Tor relays are you currently operating? Required

1-2
 3-4
 5-6
 More than 6

4. Please indicate the type(s) of relay you are running on the Tor network: Required

maximum 3 answer(s)

Guard and middle relays (also known as non-exit relays)
 Exit relay
 Bridge

Question **5** refers to the cumulative value of the 'BandwidthRate' configuration across all your relays.

Note: To ensure the highest level of participant anonymity, we want to assure you that any specific numbers provided in response to this question will be carefully anonymized. Individual identities will not be identified, and all data will be aggregated and analyzed collectively. Your privacy is of utmost importance to us, and we appreciate your trust in providing us with valuable insights.

Figure B.3: Section 1

5. What is the total bandwidth, that is allocated for all your relays combined (in GB/s)? Required

6. Could you please indicate whether the relays you are operating is up to date, or if not, how many versions behind are you? Required
- Up to date
 - One version behind
 - Two versions behind
 - Three or more versions behind
 - I don't know

7. If your relay is not up to date, please let us know the reason(s) for this, if any. Required
-
- Character limit: 0/1000

Page 1/3



Figure B.4: Section 1

Attitudes towards Current Tor Updates

Overview of the current process for updating Tor relays.

1. The Tor developers agree on the finalized source code for the update.
2. The Tor developers digitally sign the update patch using their private key, which ensures the authenticity and integrity of the patch.
3. Tor relay operators become aware that a new update for Tor is available through various sources, such as official announcements or news.
4. Tor relay operators download the update patch either from the official Tor website or by using a specific command or tool provided by Tor.
5. Before installing the update patch, Tor relay operators verify the authenticity of the patch by checking the digital signature. They ensure that the patch has been signed by the trusted Tor developers, thus confirming its validity and preventing unauthorized modifications.
6. Once the update patch's authenticity is confirmed, Tor relay operators proceed to install the patch on their relay nodes.

Figure B.5: Section 2 - Introduction

8. This question aims to assess the relative significance of different factors that may influence the decision to update Tor relays. Please indicate the level of importance you assign to each factor in influencing your decision to update: Required

	Not at all important	Low importance	Neutral	Important	Very Impor
Everything works fine, so this new update seems unnecessary	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Update takes up a significant amount of the user's time during the installation process	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Update requires unnecessary restarts	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I do not keep up with Tor update news regularly	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Update seems unimportant	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Update is related to security	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Update is related to features	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Figure B.6: Section 2

9. Continuing from the previous question, indicate the level of importance you assign to each factor in influencing your decision to update: Required

	Not at all important	Low importance	Neutral	Important	Very Impor
Updates may add unwanted features	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Updates may remove wanted features	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Updates may cause compatibility issues	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Update process may consumes too much CPU resources	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Updates may introduce new bugs	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Updates may introduce malicious content	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

10. What improvements or changes would you like to see in the update process of Tor relays? Required

Character limit: 0/1000



Figure B.7: Section 2

Tor Relay Update Attitudes and Factors Survey

Responses to this form are anonymized

Perceptions of Automatic Update Design

Comprehensive Overview of the New Update Design

The proposed update design aims to tackle challenges in the Tor network associated with protocol compatibility. Currently, multiple versions of Tor relays coexist, which poses a security risk as it may leave the network vulnerable to attacks. To mitigate this risk, the current approach is to reject relays running older software versions, even if they contribute significantly in terms of bandwidth or security features.

The new update design, called **FAN (Flexible Anonymous Network)**, addresses the complex task of maintaining a distributed network involving multiple actors, such as Tor. In the current scenario, network developers lack control, leading to heterogeneity in software versions and the need for protocol tolerance and forward-compatible strategies. To address these issues, FAN provides an architecture that is independent from OS distributions and unattended relays within the Tor network. This design enables FAN to improve anonymous communication in a more flexible way. It achieves this by facilitating on-the-fly negotiation and deployment of protocol features, customized to meet the specific requirements of different application-levels. One essential component of FAN is the **FAN Transparency Log (FTL)**, which stores critical information about updates and ensures transparency within the system.

FAN developers have three sets of actions. Firstly, they can issue and withdraw new and old updates. When issuing an update, developers send its name and attached update to the FTLs, along with meta-information containing a valid signature, protest epoch, and push epoch. The protest epoch provides relay operators with the opportunity to review the update and record any protests officially. If the protest epoch passes without withdrawal, the push epoch determines the update's availability in the Tor network.

FTLs are responsible for building the Tree for each epoch, which includes all available updates. FTLs generate proofs and respond to requests for proof of availability or absence for specific updates. Additionally, FTLs collect relay operators' signed protests updates received until the protest epoch, appending them to the update's meta-information.

Relay operators can verify, fetch, or share proofs of availability or absence from FTLs based on epoch values. They can also choose to formally protest upcoming updates before the protest epoch elapses, using their relay's identity key. These protests

Figure B.8: Section 3 - Introduction

serve as informational and globally visible indicators for FAN developers to decide whether to withdraw the update or proceed with its push epoch, as defined in the update issuance order.

The FAN design introduces software agility within the Tor network, providing developers with increased control over the update process. This, in turn, enhances the security of the Tor network.

11. Now, assuming that updates can be done automatically as described above, please indicate the level of importance you assign to each factor in influencing your decision to approve a new update patch: Required

	Not at all important	Low importance	Neutral	Important	Very Impor
Automatic updates may add unwanted features from programs	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>
Automatic updates may remove wanted features from programs	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Automatic updates may cause compatibility issues	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Automatic update process may consumes too much CPU resources	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Automatic updates may introduce new bugs	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Automatic updates may introduce malicious content	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Figure B.9: Section 3

12. Based on the automatic update design described above, where Tor developers can push updates to all relays after the protest epoch without relying on relay operators for further action, please indicate your level of agreement with the following statement: Required

Stongly disagree Disagree Neither agree nor disagree Agree Stronly

This proposed design introduces a single point of failure that could potentially impact the entire network's security

13. Based on the automatic update design described above, where Tor relay operators' only required action regarding the update is whether to officially protest the upcoming update, please indicate your level of agreement with the following statement: Required

Stongly disagree Disagree Neither agree nor disagree Agree Stronly

This proposed design reduces the level of control relay operators have over the update process, potentially affecting their ability to customize their relays

The design simplifies the update process for relay operators, making it more user-friendly

Figure B.10: Section 3

14. Based on the automatic update design described above, which provides relay operators with the opportunity (with cryptographic integrity protection) to "protest" updates before they are pushed out to the network, please indicate your level of agreement with the following statement: Required

Stongly disagree Disagree Neither agree nor disagree Agree Stronly

This proposed design actively incorporates feedback from relay operators, allowing them to have a voice in the decision-making process for updates

15. Is there any other feedback or suggestions you would like to provide regarding automatic updates in Tor? Required

Character limit: 0/1000

Figure B.11: Section 3

Appendix C

Detailed Demographic Information

Table C.1: Characteristics of Tor Relay Operators Participants

Demographic	# Participants	
<u>Length of Tor Usage for Online Privacy and Anonymity</u>		
Less than 1 year	5	
1-3 years	4	
3-5 years	8	
More than 5 years	37	
<u>Tor Relay Operator Experience</u>		
Less than 1 year	8	
1-3 years	15	
3-5 years	14	
More than 5 years	17	
<u># Operating Relays</u>		
1-2	32	
3-4	8	
5-6	3	
More than 6	11	
<u>Operating Relay Types</u>		
<i>One type only</i>	Non-exit relays only	20
	Bridge only	13
	Exit only	4
<i>Two types</i>	Excluding Exit	11
	Including Exit	2
<i>All three types</i>		4

Appendix D

Wilcoxon Signed-Rank Test Results

Table D.1: Normality Testing of Influential Attitude Factors Under Current and New Update Frameworks

Factors	Shapiro-Wilk (Sig.)	Factors	Shapiro-Wilk (Sig.)
UF1-Unwanted	<0.001	AUF1-Unwanted	<0.001
UF2-Wanted	<0.001	AUF2-Wanted	<0.001
UF3-Compatibility	<0.001	AUF3-Compatibility	<0.001
UF4-CPU	<0.001	AUF4-CPU	<0.001
UF5-Bugs	<0.001	AUF5-Bugs	<0.001
UF6-Malicious	<0.001	AUF6-Malicious	<0.001

		Ranks		
		N	Mean Rank	Sum of Ranks
AUF1-Unwanted - UF1-Unwanted	Negative Ranks	6 ^a	8.67	52.00
	Positive Ranks	19 ^b	14.37	273.00
	Ties	29 ^c		
	Total	54		
AUF2-Wanted - UF2-Wanted	Negative Ranks	9 ^d	8.00	72.00
	Positive Ranks	14 ^e	14.57	204.00
	Ties	31 ^f		
	Total	54		
AUF3-Compatibility - UF3-Compatibility	Negative Ranks	8 ^g	10.31	82.50
	Positive Ranks	15 ^h	12.90	193.50
	Ties	31 ⁱ		
	Total	54		
AUF4-CPU - UF4-CPU	Negative Ranks	7 ^j	6.50	45.50
	Positive Ranks	9 ^k	10.06	90.50
	Ties	38 ^l		
	Total	54		
AUF5-Bugs - UF5-Bugs	Negative Ranks	7 ^m	8.43	59.00
	Positive Ranks	15 ⁿ	12.93	194.00
	Ties	32 ^o		
	Total	54		
AUF6-Malicious - UF6-Malicious	Negative Ranks	4 ^p	9.00	36.00
	Positive Ranks	18 ^q	12.06	217.00
	Ties	32 ^r		
	Total	54		

- a. AUF1-Unwanted < UF1-Unwanted
- b. AUF1-Unwanted > UF1-Unwanted
- c. AUF1-Unwanted = UF1-Unwanted
- d. AUF2-Wanted < UF2-Wanted
- e. AUF2-Wanted > UF2-Wanted
- f. AUF2-Wanted = UF2-Wanted
- g. AUF3-Compatibility < UF3-Compatibility
- h. AUF3-Compatibility > UF3-Compatibility
- i. AUF3-Compatibility = UF3-Compatibility
- j. AUF4-CPU < UF4-CPU
- k. AUF4-CPU > UF4-CPU
- l. AUF4-CPU = UF4-CPU
- m. AUF5-Bugs < UF5-Bugs
- n. AUF5-Bugs > UF5-Bugs
- o. AUF5-Bugs = UF5-Bugs
- p. AUF6-Malicious < UF6-Malicious
- q. AUF6-Malicious > UF6-Malicious
- r. AUF6-Malicious = UF6-Malicious

Figure D.1: Wilcoxon Signed-Rank test ranks table

Appendix E

Themes and Coded Responses

E.1 Themes and Coded Responses of Q10

E.1.1 Lacking Update Communication

E.1.1.1 Insufficient Pre-Update Information

P1: ...clear notification prior... update would be needed

P3: More information, possibly sent in email explaining how to carry out update.

P6: There isn't a roadmap or a todo list of incoming features/bug fixes. I don't know what's planned next or what's coming next year for example.

P14: ...Being able to be notified of pending updates outside of checking the relay.

P17: I would like some better patch notes that reflect all the changes compared to a previous version, including any consequences/changes/repercussions/effects on Tor operators...

P30: ...it's (setup for automatic update) also not sanctioned by tor project and not published on official website, so others might not benefit from it.

P47: More details on the changes and possible impact.

E.1.1.2 Inadequate Post-Update Reporting

P1: clear notification ... after update would be needed

P11: ...I would prefer to know when updates are installed in case something goes wrong.

P39: Maybe, a changelog written in the user's home directory who operates the tor-daemon (e.g. one user of group debian-tor under Debian) would be nice.

E.1.2 Complex Update Process

E.1.2.1 Difficulty Accessing Updates

P9: login to unopened sites

P14: Providing an easier way to validate updates before installing...

P17: ...Sometimes it takes a while before the FreeBSD packages are updated. Would be great to speed this up if possible (especially when there are security updates).

P22: ...However, there have been times in the past where the Debian version of Tor is a few updates behind the current official version.

P37: Easy links to Debian backports to stay up to date

P41: Official repos for more linux distros...

P43: Faster auto-recognition of fact update exists...

P54: maintain debian repositories to provide updates

E.1.2.2 Preference for Automation

P1: some automation would be nice...

*P11: I would like to see the *option* of fully automated updates...*

P30: An easier setup for automated updates on “official” tor site. Currently I manually follow 5 steps, takes 20 minutes to setup...

P50: ...auto updating with properly signed / authenticated procedure would be good...

E.1.3 Tor Auto-Update is Simple

P7: ...my tor relay updates itself whenever there is an update. this is done in the way in which it is described on the tor website.

P15: ...The process is fine, works through my package manager.

P38: Installing/updating using system’s package manager (as I’m already doing)

E.1.4 Update Side Effects

E.1.4.1 Update Penalization Concerns

P12: ...sometimes the relay does not reconnect as expected...

P12: ...I feel the need to reboot and then I obtain a new IPv6 IP.

P36: Not losing the guard flags after a upgrade or restart would be nice

P48: ...shared a concern with other operators that the network's routing algorithms seemed to penalise those who updated Tor (and other parts of their systems) regularly and promote unsafe "uptime at all costs" behaviour by prioritising relays with higher uptime. I do not know whether this has been fully addressed.

E.1.4.2 Machine Stability Concerns

P34: more (smaller) updates

P40: the update process needs to be automatic, seamless, and secure. just don't ruin my machine.

P41: ...especially lightweight ones (updates)

P43: ...better insurance tor updates (or Git repos, re snowflake) aren't being hacked.

P46: Please ensure tor relay operators are never forced to accept an update, not with a time limit, not with some voting process, not for any reason...

P50: ...Yes auto updating with properly signed / authenticated procedure would be good if totally secure...

E.1.5 Update as Instructed by Tor

P50: ...Updating is very important and personal considerations of low importance. If Tor want an update I'll do it. Period...

E.2 Themes and Coded Responses of Q15

E.2.1 Support for New Auto-Updates

P37: ...Bring it on. One thing less for me to have to be concerned with.

P16: I have no problem rejecting old running software versions more quickly by the Tor...

P14: It seems important, given the centralized nature of pushed updates, that relays should be able to cryptographically verify the update before it is applied. The public signing keys need to be transparent and have a reliable method of revocation.

E.2.2 Desire to Opt-Out of Auto-Updates

P1: couldn't there be an opt-out option as well and if you don't update in time you fall off the network as today?

P10: It should be opt-in only. Or at least opt-out should be possible.

P12: If the number of bridges/relays operated is below a certain threshold, maybe still allow manual updates, if possible.

P16: ...but I would like to be able to configure an automatic update feature as an option. (Autoupdate: 'Yes' or 'No'. This is how we do it in the config of the firmware for the Freifunk routers in our community)..

E.2.3 Recommendations for Implementation

E.2.3.1 Accessibility

P15: I feel like different OS versions need to take in the updated versions? ...

P22: Not all Tor relay operators speak English. If someone makes a protest that is not in English, is there some way to translate their comment into English and vice versa if they want to read protests made in English when English is not their spoken language?...

P48: The automatic update process will need to consider network diversity - will the updates be distributed as source code to allow building for any OS/architecture or will the network restrict the ability of operators to run their choice of platform by only distributing a limited range of pre-built binaries?

P49: How would this work for operators using Tor from a package repository, where the repository controls the version?

E.2.3.2 Update Timing

P11: It really depends on the protest epoch. The last thing I need is yet another "You must review this NOW or we're taking potentially damaging action against your network without your consent." I would lose at least two relay locations where I have assured the network owners that I personally vouch for all updates.

P15: ...So it's not like they can just update everyone immediately

P16: ...(I) would like to try upgrades on test relays first and then I want to upgrade my servers gradually, one at a time. With a delay of seconds, minutes, hours or days depending on the new Tor features...

P16: ...If all servers are upgrading at the same time and there is a problem, a large part of the Tor network will be offline...

P36: It may be useful to have the option to delay non-critical updates for a limited time (like 24 hours up to 7 days), or having a limited rollout of new versions of for,

instead of upgrading the entire For network at once.

P40: ...Maybe a tiered approach to versioning; divide the network into thirds, one third is at the forefront of being updated, one third get updated every couple of months, the last third gets updated once a year and the relay operator gets to pick the tier.

P47: Possibly add criticality for updates and tie it to the period allowed for protests.

E.2.4 Opposition to the New Auto-Updates

E.2.4.1 Satisfaction with Current Updating

P7: ...My tor relay is a dedicated system only for the tor relay with system updates and tor updates periodically running and keeping everything up2date to keep the whole thing low maintainance. There where months where I didn't even log into the system to check if everything is fine. Never really had any troubles with running the system this way.

P19: I'm really not sure automatic updates are so much needed: the network works well in its current state. Relays being pushed out of the network if they are not up to date is also a way to remove abandoned relays...

P20: I would not use an update mechanism decoupled from the OS main updating process (in my case, Debian's apt)...

E.2.4.2 Concerns Over New Auto-Updates

P16: ...Most importantly, I want to update when I'm (at) home & online and able to intervene...

P17: It doesn't make sense and adds many risks to the Tor operator's infrastructure. As a Tor operator I'm responsible for and in-control of my hardware, software and network. TPO can't take-over this responsibility with some sort of managed appliance-like software. If TPO wants this, then they can just run their own hardware and be their own system administrators ;-)...

P17: ...But more fundamentally it looks like a bad idea to me. Instead of involving the operators beforehand TPO implements big changes they would just make/build/code it and then ask for a majority vote on it before effectuating the change to all relays? Sounds inefficient but also suboptimal. A healthy community has checks, balances and involvement up front instead of afterwards....

P17: ...Also I don't see what problem will be solved. The majority of outdated relays are the result of laziness/time-shortage/forgetfulness and not some deliberate

consideration or disapproving with new features or something.

P19: ...It's not only a question of Tor's update, but a question of general maintenance of the relay's servers.

P27: Approach sounds like turning relays into something like BOINC clients at best and a botnet at worst.

P30: Due to low participation in the community, it's likely no one will check the update during protest period. In that case, the system design might not improve security in reality – the developers still dictate features...

P30: ...Also, it seems to make the entire network easier to “fork” and left out a fraction of relay operators when they actually object a new feature. It's more important to not fracture the community, even when in the future some operators/developers/etc. start to disagree and can't reach a consensus.

P32: This update would seem to introduce a lot of complexity into the update process, and being unable to change compile-time parameters on unconventional build environments (custom openssl directory, prefix...) seems like a drawback as well...

P32: ...Would this change enable sybil attacks to prevent security patches from percolating the network?...

P32: ...Should relay operators to have this responsibility? It's maybe a 'feature' to know which relays are running on systems/by people who actively tend to and maintain their systems where we can assign some credence to how likely a relay is to be compromised, say, if it's likely other software on the system is outdated as well. Do we want unattended, zombie relays to continue transiting traffic if they're not running on systems which are actively maintained?

P33: I dislike automatic updates for all kinds of software, but for Tor specifically, it seems even worse.

P34: ...I'd rather update myself.

P39: I have my doubts that the proposed approach makes so much sense...

P39: ...as a simple relay operator you probably don't have the overview of fine-grained implementations of certain goals (e.g. “tamper protection”)...

P40: The proposed design process is clearly meant to disenfranchise relay operators...

P46: ...Some (relay operators) are not even technical, i.e. the tor operator might be unavailable for a week to fix any issues broken by a central authority and a working relay becomes a broken relay.

P52: This proposal has a strange smell to it.