A Survey to Guide Group Key Protocol Development

Ahren Studer, Christina Johns
Carnegie Mellon University
Dept. of Electrical & Computer Engineering
{astuder,cjohns}@ece.cmu.edu

Jaanus Kase
Carnegie Mellon University
Human Computer Interaction Inst.
jkase@andrew.cmu.edu

Kyle O'Meara
Carnegie Mellon University
Heinz School of Public Policy & Mgmt.
komeara@andrew.cmu.edu

Lorrie Cranor
Carnegie Mellon University
School of Computer Science
lorrie@cs.cmu.edu

Abstract

A large number of papers have proposed cryptographic protocols for establishing secure group communication. These protocols allow a set of group members to exchange or establish keys to encrypt and authenticate messages within the group. At the same time, individuals outside of the group cannot eavesdrop on group communication or inject messages. There have even been usability studies, demonstrating an average user can successfully complete some of these protocols. However, group protocols are rarely used in the real world. In this work, we conduct a survey to help uncover why the general population ignores such mechanisms for group communication. We also try to determine what protocols would best match respondents' current expectations for group protocols and methods for establishing trust. Survey results indicate that a group protocol that leverages location-limited channels, PKI, or Web-of-Trust authenticated public keys and allows addition and deletion of members fulfills the majority of users' expectations.

1 Introduction

Group key protocols allow a number of individuals to securely exchange cryptographic keys or establish a shared key using an insecure medium (e.g., wireless or Internet connectivity). After forming a group, members can encrypt, decrypt, and authenticate messages to and from other members of the group. Provided secure underlying cryptography, no one outside of the group can eavesdrop on the communication or inject a message that will successfully authenticate. Prior works on group protocols [2, 5, 19, 21] often use examples of collaborating researchers at a conference to motivate their work. In addition, the general population also naturally forms groups to communicate about potentially secret information. Friends try to plan surprise parties. Business partners collaborate on new projects. A group of doctors may want to discuss a specific patient's condition. However, people rarely use group key protocols to secure their communication. The goal of this work is to uncover why "average users" do not use these protocols and to determine which protocols match users' mental models.

Prior work has shown users can successfully complete various tasks associated with many group security protocols (i.e., password entry, string comparison, and data verification) [20]. The goal of this work is to determine what type of group protocols match users' paradigms and what functionality is necessary for adoption. Kuo et al. [13] analyzed group protocols with respect to different social requirements, but did not collect any end user data. Rather than postulating what users' want in a group key protocol, we use a survey to help gain insight into users' threat models and group interaction habits. To help answer the question of why users ignore group protocols, our survey was designed to help answer several sub-questions about group communication: do people not worry about protecting their communication, do current protocols not provide the necessary group management functionality, and do people only meet in scenarios where these protocols are inapplicable (e.g., a protocol that uses infrared communication [4] ceases to work when individuals communicate over the Internet). Using responses to the survey, we can propose what type of group key protocols best match current users' practices.

Our results indicate that a large portion of respondents have a very weak attacker model where attackers have limited capabilities and service providers and servers are trusted. Respondents' limited paranoia matches their practices in that only a small fraction utilize security software to provide

1063-9527/08 $25.00 © 2008 IEEE
DOI 10.1109/ACSAC.2008.28

end-to-end security for their current communication. Responses indicate that when dealing with current groups, respondents require a mechanism that allows dynamic group management (adding or removing members without having to reform the group). When meeting in person, most respondents trust third parties or physical interaction (e.g., exchanging business cards) to help verify online identities that are encountered later. Such findings indicate location-limited channels and public key based protocols correspond with the majority of respondents' trust models. When meeting online, respondents did not agree on a single solution to verify an identity. For successful adoption, online group key protocols must accommodate multiple mechanisms to establish trust or require a change in user paradigm.

The remainder of the paper is organized as follows. In Section 2, we provide an overview of previous work. In Section 3, we introduce our survey methodology. In Section 4, we present the results from our survey and answer each of the questions presented in the survey goals section (Section 3.1). After presenting the survey results, Section 5 contains guidance on what type of group key protocols would match most survey respondents' current group paradigms and habits. We make concluding remarks in Section 6.

2 Previous Work

A large number of papers have been written on the problem of group key establishment and exchange. Group key protocols allow a number of individuals to exchange or establish a shared key. Once every member of this group has a copy of the key, group members can encrypt and authenticate messages within the group. Prior works have focused on making the protocols more efficient (e.g., requiring fewer rounds of communication [1, 19]) or proving security under certain assumptions [7]. Our work focuses on the usability of the protocols and examines what group structures match users' paradigms and what functionality is required in a group key protocol. Only a limited number of other works examine the usability of group key protocols. Uzun et al. [20] examined the usability of pairing methods which overlap with a number of the tasks involved in group protocols. Kuo et al. [13] performed a survey of group protocols and analyzed how different social scenarios warrant different group mechanisms (i.e., friends sharing pictures have different expectations and threat models than an underground organization). Other works [8, 9] have examined usability associated with encrypted emails.

Group key protocols can be divided into 4 different categories based on how trust is established: public key, password, string comparison, and location-limited channel based protocols. Public key based protocols [2, 5, 11, 12, 17, 18] rely on authenticated public keys. The group members' contributions are combined such that only a member who contributed a key can calculate the shared group key. If the public keys are not authenticated, a man in the middle attack is possible [13]. Such authentication requires either a Public Key Infrastructure (PKI) or a web-of-trust mechanism (i.e., the system used in PGP [23]) to ensure that the public key one member receives corresponds to another member's actual contribution (as opposed to an attacker's contribution).

In password-based protocols, each member of the group uses a shared short secret (i.e., a human memorable password) to encrypt their messages used to generate the group key [1, 2, 7, 21]. Only a user with knowledge of the password can decrypt the messages and calculate the group's key. For this type of protocol to work the group must first distribute the password using some secure channel (e.g., a face-to-face meeting or phone call, or messages over a trusted network such as a VPN).

Comparison-based group key protocols [6, 10, 14, 21, 22] rely on detection of attacks. During the protocol any member who contributes to the key can calculate the final group key. Unlike the public key or password-based schemes, an attacker can inject and eavesdrop on all messages during a comparison-based group key protocol. However, the last stage of the protocol is a comparison of a checksum of the protocol messages. If only valid members contribute to the key generation, each member's checksum will match. If an attacker inserts a value, group members will have different checksums, detect the attack, and rerun the protocol until only group members contribute to the key calculation and the checksums agree. Here members need a secure channel for comparing checksums. When members meet in person they can talk or show each other the checksum. If members are in separate physical spaces, a trusted channel (e.g., corporate VPN, trusted server, or policy-protected phone calls or text messages) is necessary.

Location-limited channels reduce the burden on users in group security protocols [4, 15, 16]. Rather than relying on public keys, passwords, or string comparison to prevent or detect attacks, these protocols leverage channels that are infeasible for an attacker to control to communicate. These channels include infrared, visual, and physical wires. These protocols are mostly used for pairing two devices, but can accommodate groups when one device acts as a communication hub to securely pass messages to all of the other members' devices. Balfanz et al. [3] introduced instant PKIs (iPKIs) where the users uses a location-limited channel to learn about the certificate authority and acquire a certificate.

There has been only limited work on the human aspects as-

sociated with group key protocols and secure communication. Uzun et al. [20] performed a usability analysis of a number of pairing methods. In this work, they examined what kind of error rates occur in a number of tasks, such as string comparison, that are used in both pairing and group protocols. Their results show that – with the right design – users can successfully complete the necessary task 95% of the time. Garfinkel et al. [8] found that in a study environment, users will use secure email. However, a study by Gaw et al. [9] found users considered the use of encrypted communication "paranoid" and avoided its use for general communication. Kuo et al. [13] examined different group scenarios and proposed a number of guidelines and properties depending on the social situation related to a group. For example, a group of friends would consider each other peers and would prefer a leaderless group. However, businesses work in a hierarchy and expect some individual (e.g., a manager) to act as a leader in a project group. Kuo's work is closely related to our project in that both examine what functionality would best support various real-world groups. However, our work has the benefit of feedback from a large number of potential users.

3 Survey

3.1 Goals

We designed a survey to help us answer several questions about why group key protocols are not widely used and what functionality is needed to help spur adoption of these protocols. We further break this issue down into a number of subquestions: 1) do people worry about protecting their communication? 2) how do people currently manage groups? and 3) how do group members meet and establish trust? The first question helps determine why group protocols are not used. The second and third questions help define what functionality and methods, respectively, are needed to help make group protocols consistent with users' current practices.

Users' perceived threats and how users try to secure communication are important factors when evaluating whether the general population will adopt the use of group protocols. If users feel their current communication is already secure (i.e., attackers cannot access valid messages or forge new messages), a user will simply send messages in the clear rather than focusing on the secondary task of security. However, if users recognize the potential threat and want to secure communication, group protocols present a more efficient solution to security than pairwise keys.

How users manage groups defines the functionality necessary in a group key protocol both during and after formation. We need to know if generating a new group makes more sense than modifying a current group. If groups are static, a protocol that only addresses group formation is appropriate. However, if new members are often added or current members are removed, users need a mechanism to adjust the group rather than constantly forming new groups.

How group members meet is a crucial factor in determining which group protocols are applicable and how users think about establishing trust between members. If members first meet in person as a group, users can leverage location-limited channels or comparison-based protocols to establish a group key. If members first meet in person but at different times, users can exchange a password or leverage location-limited channels to securely exchange public keys for later use during password or public key based protocols, respectively. If group members never meet directly, a third party is needed to establish trust between members. A trusted authority can act as a certifying authority and sign each member's public key. If group members trust mutual friends, a web-of-trust will allow users to verify public keys.

3.2 Survey Design

To answer the questions from the previous section, we asked respondents 9 open-ended and 14 multiple choice questions using an online survey hosted on Survey Monkey[1]. We asked open-ended questions to reduce any bias and allow respondents to respond in ways we may have not considered. The survey included questions on respondents' current electronic communication and security habits, how respondents manage groups, and how respondents meet and establish trust with other individuals with which they communicate.

To determine how people currently communicate and protect that communication, we asked respondents several questions about how frequently they use different electronic communication mechanisms and how comfortable they feel with each type of communication. We also tried to gauge their perception of the secrecy of that communication – both their beliefs about how difficult it is for others to access their communication and what steps they use to protect the communication. Answers to these questions help us understand what types of communication respondents are comfortable using and respondents' perceived need for and willingness to use security techniques.

We used scenarios to determine how people currently manage groups. First, we asked respondents in which area of their life would they most likely use online communication: work, school, social situations, and personal finance. Depending on the response, we described a relevant scenario

[1] A copy of the survey is available at
http://www.ece.cmu.edu/ astuder/papers/acsac08survey.pdf

and presented questions about forming a group, adding a member to that group, and removing a member from that group. The responses to these questions help us understand what group management functionality is needed to match respondents' habits.

The last portion of the survey contained questions to help determine how respondents meet other potential group members and establish trust with them. Once we know how respondents establish trust in their online correspondents, we can infer what type of group protocol is most appropriate. We focused on two general situations: when correspondents meet in person before corresponding online, and when correspondents first meet online.

For correspondents that met in person, we wanted to determine how often the respondents meet in groups and how they later confirm the online identifier matches the person they met (e.g., verify screen-name "Jim2423" is really James Smith from the meeting last week). If respondents rarely meet in groups in person, group protocols will need to utilize something other than comparison to secure the group key (since those protocols assume the entire group is simultaneously present). However, if respondents share a password with potential group members when they meet in person, password based protocols are applicable. If group members share a trusted third party (a shared friend or authority) or exchange authentic keys using location-limited channels, public key based group key protocols correspond with respondents' current practices.

The final portion of the survey was intended to collect data about how respondents establish trust with entities that they only meet online. Given that respondents rarely exchange keys with online correspondents, we asked questions related to how they verify an online identity corresponds to a claimed physical identity. Once we know how respondents establish trust with these correspondents, we can infer what type of group protocols would make most sense for use with members who do not meet in person. Without meeting in person, a group protocol using a location-limited channel is inappropriate. However, if correspondents share a common friend or trusted authority, a web-of trust or a PKI can help exchange authentic public keys. If correspondents trust a server, the two can use secure communication with that server (e.g., TLS or SSH) to act as a secure channel for string comparison during group protocols.

3.3 Survey Methodology

In this section, we discuss how we recruited subjects for our survey and how this impacts the responses and our results. To help encourage participation, each respondent that completed the survey was entered into a raffle for one of three $50 gift certificates. Our survey was advertised from 4/25/08 to 5/16/08 on the university's electronic forum, the Craig's list (`craigslist.org`), and on one of the author's blogs. The survey itself was posted on Survey Monkey, a survey website. On average it took 14 minutes for an individual to complete the survey.

More technically experienced users browse these venues which biases our survey population towards more technically inclined individuals. With greater technical experience, our population includes a larger number of users that may consider security as an issue when communicating online. A greater interest in technology also may bias users to try new technologies, such as security software, as opposed to average computer users who are slower to adopt new software or services and focus more on completing a task.

In addition to more technical experience, a large portion of our respondents were students. 29% of respondents used `.edu` email addresses and 70% claimed to use online communication for "school." As such, our survey results include data from a disproportionate number of younger and more educated individuals.

Despite the large number of young, educated, and technologically experienced respondents, the data from our survey represents the population as a whole better than prior works on group key protocols that only considered how security researchers think individuals would manage groups [13]. We expect the population we sampled would be more interested in adopting group key protocols than the general population and thus our results are likely to be optimistic in their estimates of adoption rates.

4 Survey Results

In this section, we describe the results from the survey and interpret how users responses relate to the general issue of group key protocols. After presenting some general information about subjects' responses, we present each section of the survey in detail.

192 respondents started our survey and 150 completed the entire survey. Further analysis showed that a small number of users (5 or less) had entered possibly false information in an attempt to quickly complete the survey and gain a chance to win the prize. Specifically, these respondents answered several open ended questions with symbols (e.g., "-", ".") so the survey would let them continue and/or claimed to have never met someone in person and later communicated with them online. However, as the number of suspicious answers was small and all respondents included some non-

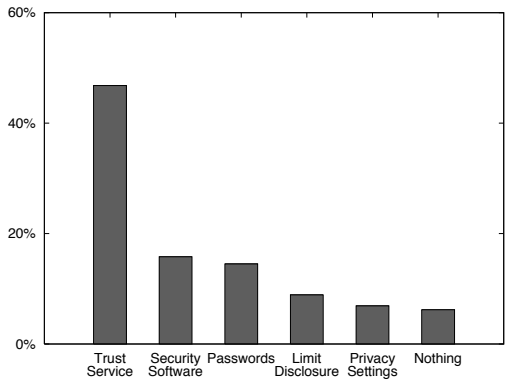

Figure 1. How respondents currently try to secure their electronic communication

suspicious answers we include all 150 completed survey responses in our final analysis.

4.1 Current Communication Habits

The first goal in our survey was to determine: how our survey respondents currently communicate online, what kind of secrecy they felt was possible when communicating online, and how they ensured secrecy. To learn how respondents communicate, we asked if they used various technologies daily, weekly, monthly, every few months, or never. Not surprisingly, respondents indicated they used email daily. On average, chat programs, email lists, social networking, and mobile messaging were used weekly. One third of the respondents reported not using group pages such as Google groups or Yahoo groups, however it is hard to determine if they considered email lists through these services as "using" the tool or just posting to a normal email list (a separate category in our survey).

We also asked respondents to rate their comfort with using these programs on a scale from 1 (unfamiliar) to 5 (very comfortable) and if they felt their communication was secure (impossible, difficult, or easy for others to access their communication). Respondents were most comfortable with email (average comfort rating of 4.88) and least comfortable with group pages (average comfort rating or 2.96). In terms of security only 3.3% felt that their online communication was impossible for others to access, 60% felt it was difficult for others to access and 36.7% felt it was simple. While people felt a high level of experience with the tools they used they still did not feel that their communication was secure, possibly because they did not think the tool can provide security.

Overall we found that respondents used chat and email the most, which are direct forms of communication that can easily adapt to group applications. Even though respondents feel experienced with these forms of communication they do not think that they are secure.

In order to see what people are currently doing to protect their communication we asked an open ended question on how they prevented unwanted access. We grouped responses into categories (see Figure 1). We found that 47% of respondents trusted a service provider (i.e., look for https to webmail), 16% of respondents used security software (e.g., PGP, encryption in Skype, or chat plug-ins), 14% relied on passwords, 9% limited disclosure (i.e., only discussing private information in person), 6% used privacy settings, and 7% said they take no steps to protect their communication.

The largest step taken to protect data is to trust services to deliver it to intended recipients. Despite knowing others might be able to see their communication very few respondents invested the time to use secure software. In fact one respondent said he "previously used Off-the-record, an encryption plugin available for some IM clients, including Pidgin. Stopped when I realized that for me personally, it wasn't worth the inconveniences." Group key protocols have a difficult road towards widespread adoption since frequent and experienced users of electronic communication often simply leave messages in the clear despite recognizing a threat to their privacy.

4.2 Current Group Management

To collect information about how respondents currently manage groups we presented each individual with a group communication related scenario that had a good chance of matching something they had done in real life. Responses to these questions help indicate what type of paradigm makes sense during group formation, and how dynamic real life groups are. Responses indicate respondents already utilize face-to-face and multicast mechanism to setup groups so group protocols do match their habits. In addition, respondents recognize groups are dynamic and depending on the situation prefer to modify the group, rather than having to generate a new group when membership changes.

To ensure the task was applicable to a respondent, we asked the respondent what type of electronic communication was most common to them: business, social, managing medical records, and personal finances. Of the 150 respondents to complete the survey, 145 felt their communication was at risk so only those respondents were asked how they managed groups. Of the 145, 87 indicated their electronic communication most often involved social situations. These 87 respondents were presented with a group scenario related to

organizing a surprise party. The other 58 were more accustomed to communicating about business and were presented with a scenario related to managing a business project. No respondents indicated their online communication was related to the management of medical records or finances.

4.2.1 Group Formation

The first task for any group is to define who is in the group, inviting the proper individuals, and ensuring only they can receive and send the relevant messages. The first question was an open-ended question about how respondents would first invite people to the surprise party or organize the first meeting for the business project. We grouped responses into the following categories based on what communication mechanism a respondent would use.

1. *Invalid answer.* These responses were off-topic or only indicated a limited grasp of how to protect the information (i.e., "do not tell the guest of honor" in the surprise party scenario).
2. *Insecure unicast.* These responses included the use of instant messages or other one-on-one mechanisms that are in the clear.
3. *Insecure multicast.* These responses included generating an event page on Facebook or sending an email to the group in the clear.
4. *Secure unicast.* Theses responses included the use of phone lines or SMS messages that provide policy protected one-on-one communication channels.
5. *Secure multicast.* These responses included encrypted emails, corporate wide secure networks/VPNs, and secured corporate internal Wiki pages.
6. *Face-to-face.* Some respondents felt the most secure method with the least likelihood of having communication leaked was simply meeting in person.

These are useful categories since they relate to what group key protocols would be applicable. With an insecure multicast and secure unicast, respondents could transmit a password, compare strings, or emulate a location-limited channel using the secure unicast and performing all other group key protocol messages in the insecure multicast medium. If respondents have a secure multicast channel, they already have some trusted public key to identify the server or the VPN and could use that as an authority to verify other public keys in a group key protocol. If respondents meet face-to-face, most group key protocols – except for public key protocols – could work as the authors originally intended.

Figure 2 indicates what percentage of respondents used each communication mechanism in each of the scenarios. In both scenarios, a small fraction of respondents ($10 - 20\%$) met in person to form groups. For a less formal surprise party, a large number of respondents suggested emailing the guests or using a Facebook event (an insecure multicast) with a significant fraction using SMS messages or phone calls (a secure unicast). If respondents were to combine both of these methods, respondents could securely form groups. However, it is a difficult task to convince respondents to do twice as much work to secure a group when little is done to secure one-on-one communication (see Section 4.1). For the business scenario, insecure email lists were common, but a significant fraction mentioned using encrypted emails and secured corporate networks and servers. If respondents are willing to trust company servers to protect communication, the respondents will also trust the servers to provide certified public keys for other members in their business project group, allowing for simple use of public key based group key protocols. The reason for the discrepancy may be the importance and duration of the groups. Business groups last for a long time, and corporate espionage is a real issue, requiring companies to protect their communication. Surprise parties are rarely planned long in advance and the guest of honor rarely acts as an adversary trying to access invitees' communication. Social groups may see security as irrelevant with such a small window of opportunity and little to lose.

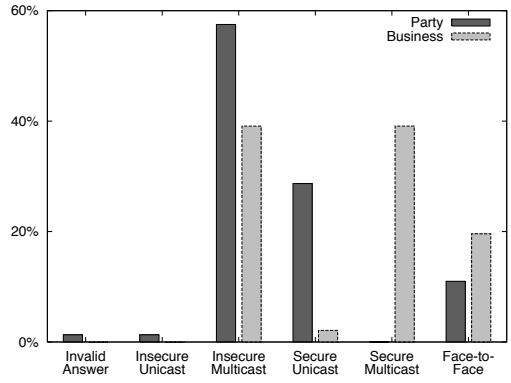

Figure 2. How respondents currently set up groups

4.2.2 Adding Members

Once a group is formed, additional members may want to join the group. To determine how respondents would handle such a situation, we asked respondents how they would invite a guest they initially forgot about to the surprise party or how they would add a new employee to the project. 78 of 87 (90%) party planners and 49 of 57 (86%) project managers indicated they would want to add the member to the existing group (e.g., "include them on any future party emails" or "add the employee to the server's access control list"). Only

one project manager indicated that forming a new group was an option.[2] This is strong evidence that a group key protocol must at least allow the addition of members after group formation.

4.2.3 Removing Members

In some situations, a member of the group should no longer receive communication from other group members. In that scenario, the unwanted member must be removed from the group or a new group is formed to maintain secrecy of any future group communication. To analyze how respondents would manage this task, respondents were asked how they would deal with a project member who was fired or a friend who leaked information about the surprise party. In both tasks the respondents were told the majority of the group still needed to communicate; project members needed to complete the project and guests needed to schedule a new surprise party without letting the evicted member access the information. An open-ended response was used to collect respondents answers. We grouped responses depending on whether a new group should be formed without the unwanted individual or if the old group should simply evict the unwanted individual.

The responses were heavily situation dependent. In the surprise party scenario, only 9% of respondents wanted to reconfigure the old group. 77% of party organizers preferred to simply make a new group (i.e., "setup a new Facebook event," "generate a new email list," etc.). The remaining 14% provided unclear responses. In the corporate scenario, the opposite was true. 76% of respondents mentioned changing access to resources to stop the expelled employee from accessing project information (i.e., "delete the fired employee's account," "remove privileges from the Wiki-page," etc.). Only 16% of respondents wanted to form a new group (i.e., "change the password on the wiki" which could mean a new group if a single password is used). These results indicate the protocol should accommodate dynamic groups where members are both added and deleted. However, the protocol should also be efficient so that if respondents choose to reform a new group whenever a member is removed, the overhead is not an annoyance to the respondent.

4.3 Making New Acquaintances in Person

To help understand how respondents establish trust and how often groups meet, we asked a number of questions about communication habits when first contact is made offline and in person. Based on these questions we can infer how respondents begin to trust online identifiers and how often groups do meet in person. Results indicate that as expected, people often meet in person before they talk online. In addition, groups meet in person a significant fraction of the time and often share a trusted third party that could facilitate the exchange of authentic public keys. Such results indicate that when respondents meet in person a wide range of group key protocols agree with current respondents' practices.

To collect information about how respondents meet and later communicate, we asked our respondents the following initial questions.

- Do you meet people in person, and then communicate with them online afterwards? (Add them to your IM list, Add them as a Friend on Facebook, etc. Send them an email.)
- After meeting new people in a group situation, how often do you only contact one person versus contact several people from a group situation?

96% of our respondents answered "yes" to the first question. This is not surprising, as our survey recruitment was done largely over the Internet, and we did expect our respondents to use the Internet to communicate. We are not sure how to account for the remaining 4% – perhaps it is survey noise or they only communicate in real-life with their real-life friends and have a separate category for online friends, and there is no overlap between these two categories.

To determine how often users meet in groups we asked the following question.

- After meeting new people in a group situation, how often do you only contact one person versus contact several people from the group situation? (1 = "100% of the time only contact an individual", 3 = "50% of the time only contact an individual, 50% of the time contact multiple people", 5 = "100% of the time contact multiple people", avg = 2.44, median = 3, n = 142)

85% of respondents checked some option other than 1 (implying a large portion of respondents meet and later communicate with groups). This result shows that groups physically meet and thus location-limited and comparison protocols are applicable.

Following this were questions that asked users to rank how often respondents met someone with shared affiliations and friends of friends.

- When you first met these people, how often did you already have a friend (not just an acquaintance) in common? (1 = never, 5 = always, avg = 3.33, n = 142)
- Do you often meet these people through a shared school, business, or other organization? (1 = never, 5 = always, avg = 3.37, n = 142)

[2]The remaining respondents did not indicate how the addition would be performed (i.e. "based on the person I'm inviting" or that no real group communication was used (i.e., "things must be handled carefully").

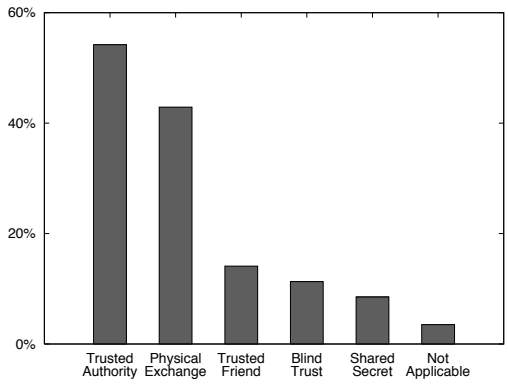

Figure 3. How respondents confirm they are communicating with the correct online identifier for a person they met

The results of these questions suggest that there is an above-average portion of respondents who meet with people offline with who they already share either a common friend or an organizational affiliation. This was further validated in the next question, where we asked about validating the contact's online identity:

- When you first communicated with one of these people online, how did you know you had found the right email address, screen name, or profile for the person/people you met?

This was an open-ended question where we partitioned the responses into six categories with the following distribution: performing background research such as a web search or directory lookup to acquire information from a trusted authority (54%), getting contact info directly from the person either verbally or in the form of artifacts such as a business card (43%), checking with a friend that the email/screen name is correct (14%), blindly trusting the person online (11%), making sure that the contact is in possession of information that the respondent, but not the general public, are aware of (8%), and not applicable (i.e., "n/a", "i dunno", "?", ...) (4%). Similar to the previous open-ended question, respondents could (and many of them did) give multiple responses, which counted as multiple categories. A comparison of these categories is shown in Figure 3. These findings support the use of public key protocols where participants could leverage common friends or businesses to verify a public key (59% of the time). However, 23% of respondents preferred a physical exchange and shared no trusted authority or friend. In such a scenario, a group protocol that utilizes a location limited channel or comparison is appropriate.

Responses in this section of the survey indicate that a large number of respondents do meet in person and often do meet with a group. In addition, these respondents often share common trusted entities with the other individuals with which they communicate. Based on these facts, the public key, location limited channel, and comparison based group key protocols discussed in Section 2 coincide with users' current habits and paradigms.

4.4 Making New Acquaintances Online

After learning how respondents verified an online identity after meeting in person, we asked questions about scenarios where respondents first met people online (as opposed to met in person). Without the ability to physically meet people some other mechanism is needed to establish trust between individuals or groups.

This section had the following questions:

- If you did not first meet someone in person, how were you first introduced to them?
  - other person contacted you
  - you found the other person through a search or browsing pages
  - introduced by a mutual friend
  - met in an online community
  - Other
- When you first communicated with one of these people online, how did you know you had found the right email address, screen name, or profile for the claimed person/people?

This first question is an important issue when it comes to establishing trust. If respondents met in a scenario with an obvious authority, that authority becomes a trusted third party for the verification of public keys. However, if a respondent simply gets an email from someone for what appears to be no reason, a different mechanism is needed to verify the individual's claimed identity. Figure 4 shows the categories and percentage of respondents that responded to each category. Respondents could mark any number of scenarios and many did so percentages add to more than 100%. Based on these responses, it is hard to infer what type of group protocol would work best for groups that do not meet in person. A large portion of respondents have received unsolicited emails where there is no obvious source of trust. In the three other scenarios (searches, friends, and communities), the web server, mutual friend, or community manager could verify public keys. The "other" category was often used, but many respondents failed to provide any answer to clarify what the other mechanism was.

To help gain a better understanding of how respondents established trust with online entities, we asked respondents an open-ended question about how they verified an online identity when someone initiated contact with them online.

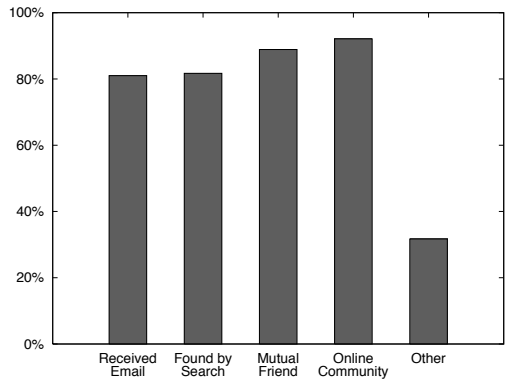

Figure 4. How respondents met people online

Respondents provided answers that we could group into one or more of seven possible categories (note percentages fail to add to 100%).

- Never trusted the identity & limited information disclosure. (28%)
- Trusted an authority to verify the identity (i.e., email address matches name or Facebook details match claimed identity). (28%)
- Assumed the person was telling the truth. (24%)
- Trusted a mutual friend to verify the identity. (15%)
- Verified the identity knew some private information – a shared secret. (13%)
- Verified the identity offline using a phone call or a meeting in person. (8%)
- Responses that were not applicable. (6%)

Figure **??** contains a summary of the results. 38% of respondents utilized a trusted friend or a trusted authority to verify an identity (scenarios where public keys are applicable). If we were to discard the more cautious users which limited disclosure, 53% of respondents could benefit from public keys (shared authority or offline verification). As such there is no clear solution that allows a large portion of the respondents to verify an online identity. Unless online communities and social networks increase in popularity and begin to act as certificate authorities (CAs), numerous mechanisms are needed to accommodate the different methods that respondents may utilize to establish trust when meeting online.

When translated to group key protocols, no real secure solution exists that would work for the majority of respondents. Only a small percentage of respondents would utilize password based protocols (13%). Location-limited channels that leverage physical collocation still are relevant since

Protocol	Management Functionality		Operates Online (w/o a Trusted Channel)
	Addition w/o Reforming	Deletion w/o Reforming	
Password Based	✔		
Comparison Based	✔		
Public Key Based	✔	✔	✔
Location Limited Channel Based	✔	✔	

Table 1. How Various Group Key Protocols Fulfill Users' Needs

8% of respondents waited for a physical meeting or an out-of-band channel to verify an identity. Public key based protocols work with the 38% of respondents that shared friends or trusted authorities that could act as certificate authorities (CAs). Provided the trusted entity is online and cooperative, these same respondents could leverage the trusted entity as a secure channel to utilize comparison based protocols. A time of first use style of authentication (similar to SSH) would allow those who blindly trust values (28%) to form groups. However, such an approach is vulnerable to active attackers that may modify or inject messages during the group key protocol. Without a single protocol that accommodates multiple methods to establish trust, multiple situation-dependent group key protocols are needed to form groups that meet online.

5 Guidelines for Future Group Protocols

Based on our findings, users will avoid adopting group key protocols unless respondents' views on security tools change because most respondents (47%) rely on service providers for security while only 15% of respondents used security software for one-on-one communications. With over 85% of respondents wanting to add members to the group and 76% of long term groups requiring deletion without reforming the group, it is clear group protocols must support dynamic groups. For groups that meet in person, verification via third parties and physical collocation agree with 59% and 43% of respondents respectively as an acceptable basis of trust. For groups that meet online, there is no single overwhelming method for establishing trust.

Table 1 compares the general types of group key protocols with respect to their management functionality and applicability to survey respondents' methods of establishing trust. Users often want to add and remove members from the group without having to reform the group(see Section 4.2). As such, public key based [12, 18] and location limited channel based [4, 16] protocols provide the necessary group management functionality. These two types of protocols agree with a large portion of respondents' trust habits when meeting groups in person (82% utilized either

method). In addition, users could utilize location limited channels to exchange public keys. With such an approach a single group key protocol would fulfill users' needs independent of how they acquired the public keys. When meeting online there is no single protocol that satisfies a large portion of respondents trust establishment habits. For successful group key protocol adoption, respondents' online practices will have to adapt to accept public key based protocols (location-limited channels do not work online) or a new type of group key protocol is needed to allow for member addition and deletion without the use of public keys.

6 Conclusions

In this work, we performed an online survey to help understand how users communicate in groups and why users continue to ignore group key protocols even though users often communicate as groups. We also investigated how users establish trust with individuals they meet both in person and online. The results of the survey indicate that users rarely secure one-on-one communication and thus securing group communication is not a goal. If users do begin to focus on security, group key protocols match users current group communication habits. For groups that meet in person, public key or location limited channel based protocols are the most appropriate. These protocols also allow the addition and deletion of members without reforming the group. However, as more communication moves online, groups that only communicate online are relevant. For online groups, respondents did not converge on a single appropriate trust establishment mechanism for group key protocols. As such, one potential direction for future work is to develop new online group key paradigms that are applicable and accepted by the majority of users.

References

[1] M. Abdalla, E. Bresson, O. Chevassut, and D. Pointcheval. Password-based group key exchange in a constant number of rounds. In *Public Key Cryptography (PKC)*, pages 427–442, 2006.

[2] N. Asokan and P. Ginzboorg. Key-agreement in ad-hoc networks. *Computer Communications*, 23(17):1627–1637, Nov. 2000.

[3] D. Balfanz, G. Durfee, and D. Smetters. Making the impossible easy: Usable PKI. In *Security and Usability: Designing Secure Systems that People Can Use*, pages 319–334. O'Reilly, Sebastopol, CA, 2005.

[4] D. Balfanz, D. Smetters, P. Stewart, and H. Wong. Talking to strangers: Authentication in adhoc wireless networks. Feb. 2002. Network and Distributed Systems Security (NDSS).

[5] M. Burmester and Y. Desmedt. Efficient and secure conference key distribution. In *Security Protocols—International Workshop*, 1997.

[6] M. Cagalj, S. Capkun, and J.-P. Hubaux. Key agreement in peer-to-peer wireless networks. *IEEE (Special Issue on Cryptography)*, 94:467–478, 2006.

[7] R. Dutta and R. Barua. Password-based encrypted group key agreement. *International Journal of Network Security*, 3:23–34, 2006.

[8] S. L. Garfinkel and R. C. Miller. Johnny 2: a user test of key continuity management with s/mime and outlook express. In *SOUPS '05: Proceedings of the 2005 symposium on Usable privacy and security*, pages 13–24, 2005.

[9] S. Gaw, E. W. Felten, and P. Fernandez-Kelly. Secrecy, flagging, and paranoia: adoption criteria in encrypted email. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 591–600, 2006.

[10] M. T. Goodrich, M. Sirivianos, J. Solis, G. Tsudik, and E. Uzun. Loud and clear: Human-verifiable authentication based on audio. In *International Conference on Distributed Computing (ICDCS)*, page 10, 2006.

[11] M. Just and S. Vaudenay. Authenticated multi-party key agreement. In *96*, volume 1163, pages 36–49, 1996.

[12] Y. Kim, A. Perrig, and G. Tsudik. Simple and fault-tolerant key agreement for dynamic collaborative groups. In *Proceedings of the 7th ACM Conference on Computer and Communications Security*, pages 235–244. ACM Press, Nov. 2000.

[13] C. Kuo, A. Studer, and A. Perrig. Mind your manners: Socially appropriate wireless key establishment for groups. *Proceedings of First ACM Conference on Wireless Network Security (WiSec '08)*, Mar. 2008.

[14] S. Laur and K. Nyberg. Efficient mutual data authentication using manually authenticated strings. In *Cryptology and Network Security (CANS)*, pages 90–107, 2006.

[15] J. M. McCune, A. Perrig, and M. K. Reiter. Seeing-is-believing: Using camera phones for human-verifiable authentication. In *Proceedings of the IEEE Symposium on Security and Privacy*, May 2005.

[16] F. Stajano and R. J. Anderson. The resurrecting duckling: Security issues for ad-hoc wireless networks. In *Security Protocols Workshop*, pages 172–194, 1999.

[17] D. Steer, L. Strawczynski, W. Diffie, and M. Wiener. A secure audio teleconference system. In *88*, volume 403, pages 520–528, 1990.

[18] M. Steiner, G. Tsudik, and M. Waidner. Key agreement in dynamic peer groups. 11(8):769–780, Aug. 2000.

[19] W. Tzeng and Z. Tzeng. Round-efficient conference-key agreement protocols with provable security. In *2000*, volume 1976, pages 614–628, 2000.

[20] E. Uzun, K. Karvonen, and N. Asokan. Usability analysis of secure pairing methods. In *Usable Security (USEC)*, Feb. 2007.

[21] J. Valkonen, N. Asokan, and K. Nyberg. Ad hoc security associations for groups. In *Security and Privacy in Ad-Hoc and Sensor Networks (ESAS)*, pages 150–164, 2006.

[22] S. Vaudenay. Secure communications over insecure channels based on short authenticated strings. In *Advances in Cryptology*, pages 309–326, 2005.

[23] P. R. Zimmermann. *The Official PGP User's Guide*. MIT Press, Cambridge, MA, USA, 1995. ISBN 0-262-74017-6.