

Improving User Decisions about Opening Potentially Dangerous Attachments in Email Clients

Ricardo Villamarín-Salomón
and José Carlos Brustoloni
Department of Computer Science
University of Pittsburgh
210 S. Bouquet St., room 6135
Pittsburgh, PA 15260
{rvillsal,jcb}@cs.pitt.edu

Matthew DeSantis
Software Engineering Institute
Carnegie Mellon University,
5000 Forbes Avenue
Pittsburgh, PA 15213
mdesanti@cert.org

Ashley Brooks
Heinz School of Public Policy &
Management
Carnegie Mellon University
5000 Forbes Ave
Pittsburgh, PA 15213
adbrooms@andrew.cmu.edu

ABSTRACT

To prevent users from opening potentially dangerous attachments, current email clients rely on mechanisms like anti-virus scanning, spam filtering, junk mail heuristics, and completely disallowing attachments of certain types. However, none of those mechanisms can guarantee complete protection against new or targeted attacks. We propose Context Sensitive Guidance (CSG), whereby the email client detects that the user is about to open a potentially dangerous attachment, interrogates the user about the context in which this is happening, and provides guidance on how to proceed. In addition, we propose Context Sensitive Guidance with Accountability (CSG+A), which also informs the user that his or her answers and decisions will be recorded, could be audited, and if found unreasonable could result in mandatory training or other penalties. We performed a pilot study to test these techniques. The pilot study achieved results that are encouraging but not statistically significant, and suggests improvements for further study.

1. INTRODUCTION

Each day, corporations have to analyze and filter volumes of illegitimate electronic mail. Virus scanners and filters have become necessary background utilities for furthering system survivability. However, anti-virus software may not protect users of very recent threats, and junk email filters may not block all new attacks. Furthermore, the risk associated with an attachment cannot be determined simply from its type: the risk depends also on context. In some situations, completely disallowing certain attachment types would hurt usability more than improve security. Although system administrators may install various mechanisms to protect users from email-borne attacks, ultimately an organization's security will still depend on users' making sound decisions on how to handle email attachments.

Email decision-making skills could be taught. However, many email users do not have such training and would not voluntarily obtain it. A quicker solution might be to improve user interfaces, such that email client software helps untrained users make more secure decisions. However, insufficient results exist to support this hypothesis.

This paper addresses three usability questions related to the security of email clients. First, given a group of nontechnical computer users, how likely is it that they will open a dangerous attachment while performing their daily activities? Second, is it possible to provide *effective* guidance to users so that they can make better decisions about opening potentially dangerous email

attachments? Third, does an audit trail enhance guidance in helping users make disciplined decisions about opening potentially dangerous attachments?

We propose two novel user interface techniques, CSG (Context Sensitive Guidance) and CSG+A (Context Sensitive Guidance with Accountability). CSG attempts to detect when the user is about to open a potentially dangerous attachment, interrogate the context in which this is happening, and provide the user specific guidance on how to proceed. CSG+A adds to CSG an audit trail. CSG+A informs the user that the user's answers and decision will be recorded and could be audited. If auditors find the user's decision-making unsound, they may impose penalties such as suspending the user's email privileges until the user passes appropriate training.

To evaluate our techniques and answer the aforementioned questions, we implemented CSG and CSG+A in the popular Mozilla[®] Thunderbird[™] email client. We performed a pilot study with three independent groups respectively comprising 5, 5, or 6 users and employing the email client unmodified, with CSG, or with CSG+A.

With an unmodified email client, we found that the incidence of users with a high likelihood of opening infected attachments was quite high. However, the incidence of users with low likelihood of opening such attachments was also high. On average, the likelihood of opening infected attachments was lower among CSG and CSG+A users. However, the differences observed between groups were not statistically significant in this small study.

The rest of this paper is organized as follows: Section 2 describes our evaluation methodology, Section 3 presents our results, and Section 4 discusses future work.

2. METHODOLOGY

The subjects of our pilot study were 16 male and female CMU students from diverse programs not related to Computer Science or Electrical and Computer Engineering. All subjects were previous users of email clients (e.g., Outlook or Thunderbird) and word-processing editors (e.g., Word). However, technical background was not a requirement. Subjects were randomly assigned to one of the three groups. The first group used an unmodified Thunderbird email client, the second group used Thunderbird with CSG, and the third group used Thunderbird with CSG+A.

Table 1. Results of pilot study

	No Guidance		Context Sensitive Guidance		Context Sensitive Guidance + Accountability	
	Total	Avg.	Total	Avg.	Total	Avg.
Attachments Opened	69	13.80	63	12.60	84	14.00
Infected attachments opened	23	4.60	17	3.40	24	4.00
Infected attachments not opened	29	5.80	35	7.00	35	5.83
Non infected attachments not opened	7	1.40	7	1.40	7	1.17

We asked subjects to role play “Chris Moore,” an employee of “ACME Corporation,” a provider of financial services to other institutions. Chris was supposed to work in a group with three people with specified characteristics. We registered the domain acmecorp.biz to create the respective corporate email addresses. We initialized Chris’ corporate email inbox with 21 unread messages, 11 of which legitimate and 10 infected. We also described to subjects some details of Chris’ private life. Chris was supposed to have a personal email account for private messages.

We asked subjects to read Chris’ corporate email and use it to complete three tasks by the end of the (one-hour long) experiment. The first task was to verify if the minutes of his group’s last meeting, as recorded by his secretary, Sally, included the decision to hire an extra worker for Chris’ group. The second task was to review job applications, select a candidate, and send Sally a request to schedule an interview. An ad for the position had been posted in job-search web sites, asking candidates to email resumés to Chris. The third task was to receive from his co-workers sections of a draft and send them a combined draft.

3. RESULTS

Table 1 presents a summary of the participants’ performance. There was a reduction in the number of infected attachments opened in the CSG and CSG+A groups (by 26% and 10%, respectively) in comparison with the control group, but the results were not statistically significant.

In large part, the inconclusiveness of the results may be due to small sample sizes and inappropriate subject screening. Our hypothesis is that CSG and CSG+A help *untrained* users make more secure decisions. However, results for the control group suggest that many users (at least those drawn from the CMU student population) already know what attachments to avoid, i.e., are not really untrained. Among well-trained users, there is little room for the user interface to have an impact.

Our experiments may have suffered from instruction bias. Subjects may have weighed more heavily our instructions that they *read* Chris’ email and *complete* the specified tasks than the software’s guidance. In retrospect, we also find that our scenario has confounding factors that could be causing cognitive overload. Several of the scenario’s infected messages have the same subject or sender as legitimate messages that we specifically instructed subjects to expect for performing the assigned tasks. Arguably, such coincidences are unlikely in practice and should not be such a predominant part of the assumed scenario. As it is,

the scenario requires subjects to make many decisions that subtly take into account a plethora of assumptions. In practice, users could be expected to be familiar with those details, but the assumed scenario may not be giving subjects enough time and opportunity to take in all the relevant information before making decisions.

4. FUTURE WORK

We are planning a new user study intended to rectify the problems encountered in the pilot study.

Pilot results suggest that the decision-making skills of all potential subjects need to be assessed with the unmodified email client, so as to (1) estimate what fraction of the population could benefit from training or a better user interface (pilot results suggest maybe half of the general population), and (2) interpret a subject’s results with CSG or CSG+A in light of that subject’s results with the unmodified email client (pilot results suggest that only if the latter are poor will CSG or CSG+A benefits be significant).

A straightforward way to achieve such calibration is possibly to have a same user group perform similar scenarios with the unmodified email client and then with a modified email client. (Given that all subjects already have used an unmodified email client, such an order does not introduce bias). Because a same subject will perform multiple scenarios, the latter have to be correspondingly simplified. Additionally, the new scenarios need to avoid instruction bias and confounding factors as much as possible.

Another avenue for improving usable security is to design novel mechanisms that enhance security but are transparent to users. Our current implementation detects that an attachment is potentially dangerous based solely on the attachment’s extension, and then interrogates the user about the context. This simple heuristic could engage users in perhaps too many dialogs, leading to user irritation or neglect. We plan to investigate also more sophisticated heuristics that can automatically determine context without asking the user, and can in some cases make transparent security decisions.

5. REFERENCES

[1] Villamarín-Salomón, R., Brustoloni, J., DeSantis, M. and Brooks, A.: Improving User Decisions about Opening Potentially Dangerous Attachments in Email Clients. Department of Computer Science, University of Pittsburgh, TR-06-136.