The Personal Information Security Assistant

Roeland H.P. Kegel
University of Twente, Netherlands
r.h.p.kegel@utwente.nl

***Abstract*—The human element is often found to be the weakest link in the information security chain. The Personal Information Security Assistant project aims to address this by improving the privacy and security awareness of end-users and by aligning the user's personal IT environment to the user's security requirements. It does this by elicitation of a user's privacy and security requirements (risk appetite) as well as a user's risk perception. The PISA then takes action by aligning the user's requirements and perceptions, thereby improving user awareness regarding privacy and security. This article outlines the research questions, methodology and current results associated with the PISA project.**

***Index Terms*—Persuasive Technology, Security, Privacy, Risk Perception, Risk Appetite, User Requirements Elicitation.**

I. Introduction

As digital systems and services become more complex and connected, they become both more vulnerable and harder to understand. As a consequence, it becomes more difficult to grasp how our actions affect our privacy and the security of our systems. After technical mitigations have been implemented, the human element often turns out to be the weakest link in the information security chain. This means that hardening the user to security and privacy threats has become of paramount importance. To do so, two factors are important to consider: a user's risk appetite (which forms the user's privacy and security requirements) and user awareness of to his or her security and privacy. We face challenges in both of these areas: to model a user's risk appetite, a wide range of constantly changing threats have to be considered. Additionally, users themselves (and their associated risk appetites) are diverse. Modeling a user's risk appetite is a significant research challenge. On the other hand, risk perceptions in information security are strongly influenced by incidents [1], [2]. Since incidents are rare, a mismatch often exists between a user's stated risk appetite and his or her risk perceptions.

The (re-)alignment of risk appetite and perception is of vital importance in order to effect any change in a user's behavior. To do this, we have to model both risk appetite and risk perception. In the Personal Information Security Assistant project[1], we aim to reduce privacy and security risks to end-users by educating them and keeping them aware of the current risk landscape, by monitoring and eliciting their current risk appetite, and by taking action to align the current risk landscape to their current risk appetite. In this sense, the PISA system acts as a kind of personalised security operations center for the individual: informing, protecting and educating its users.

To create a such a PISA system, we need to overcome several obstacles. We need to create and maintain a model of the user's risk appetite and of the risk landscape that the user is exposed to, and we need to find effective ways to align the user's behavior with his or her risk appetite. In addition, we also need to consider the risk that the PISA itself poses to security: the PISA needs to do all this in a way that does not pose a threat to privacy in itself. Our initial solution idea is to take as much as we can from currently established enterprise risk assessment methods and adapt them to the situation where risks of a single individual need to be assessed and mitigated. To inform the user and to change his or her behavior, we will also need to incorporate visualization and persuasion techniques, so we need to look into these as well. The next section considers these problems by distilling them into a set of research questions to be answered.

II. Research Questions

The primary goal of the PISA is simple: enhancing a user's privacy and security. To do so, several questions need to be answered. An associated main research question and subquestions are posed below. The main question associated with this goal is decomposed into three subquestions. First, question 1 studies the problem of user modeling: how can we characterise users and their individual requirements? Question 2 then asks for the different elements of the PISA system itself. For each of the subquestions of question 2, empirical lab and field tests are planned to test our answers. Finally, question 3 studies the risks associated with the tool itself. The methodological approach for answering these questions is discussed in section V.

Main Research Question: How can we improve how users think, act and feel regarding digital privacy and security?

Question 1: How can we design a tool that motivates and persuades users with regards to digital privacy and security?

- *Q1.1:* How can enterprise risk assessment methodologies be adapted for and applied by end-users to enhance the security of the IT systems they use, and/or their own privacy?
- *Q1.2:* How can user ability and motivation regarding privacy be measured by observing human computer interactions?

[1] http://scs.ewi.utwente.nl/projects/pisa/

978-1-4673-6905-3/15/$31.00 © 2015 IEEE

- *Q1.3:* How can a user's privacy requirements and status be visualised?
- *Q1.4:* How can the human-computer dialogue between PISA and the user be structured in a persuasive manner?

Question 2: What information is needed to achieve these changes in user behavior and attitude?

- *Q2.1:* How can we discover what privacy requirements exist for end users, and how can we determine which of these apply to specific users and/or contexts?
- *Q2.2:* What user-characterising factors can and should influence how PISA takes action and informs the user, and how do these factors relate to user requirements?

Question 3: How do we prevent the PISA from becoming a security risk in its own right?

- *Q3.1:* How can we assure the privacy of a user's data while allowing for analysis at a centralised location?

III. Challenges

We have identified the following challenges that the PISA has to overcome in order to achieve its goals:

- **Multidisciplinary:** the PISA applies existing research from a wide array of fields such as computer security, the social sciences, machine learning and information systems. Although this allows for a great amount of synergy, it also poses a significant challenge since such a wide array of literature and expertise has to be consulted.
- **Dynamic threats:** since the goal of the PISA is to protect a user's privacy and security, it not only faces the significant amount of differences that exist among people, but also has to take into acocunt the constantly changing and diverse array of threats that it needs to address. Any answer to the main research question posed in the PISA project has to take this dynamic nature into account.
- **Low Motivation & Skill:** the Elaboration Likelihood Model (ELM) [3] defines the constructs *user ability* and *motivation*, which influence what type of communication with the user will be effective. Since risk perceptions are strongly tied to the occurrence of incidents [2], we assume that both ability and motivation to improve IT security and privacy are low among the general public until a security incident occurs. The PISA aims to prevent such incidents before they happen, and so the PISA needs to employ persuasive techniques that are effective when ability and motivation are (still) low: techniques that require little intervention or thought from the user. Such a communication strategy relies on the ELM's *peripheral route*, where users use heuristics and peripheral cues to make their decisions. However, directly conflicting with this strategy is the type of change that PISA needs to achieve: sustained behavioral and attitudinal change is needed to raise security awareness [1]. This type of change is best realised through education and intervention. Such a strategy corresponds to the ELM's *central route*, relying on careful thought and consideration from the user. Any implementation of the PISA framework will require careful consideration of the human-computer dialogue to retain the user's attention.
- **Requires extensive tool infrastructure to test:** the broad nature of the domain in which PISA will operate means that, to answer the research questions posed in the previous section, a large amount of infrastructure has to be built: programs for every device and threat, separate sensors for each user property we want to monitor, as well as their associated rules need to be defined to fully test the PISA concept. Identifying the minimal amount of components needed to answer each research question poses a significant challenge.
- **Expert domain knowledge:** the mitigations to threats the PISA has to address often involve a large amount of domain specific knowledge. For every threat PISA has to defend against, expert domain knowledge is needed.
- **Ensuring privacy by design:** as stated in research question 3.1, a large challenge is preventing the PISA from becoming a security risk of its own. Since the PISA is an in-depth profiling tool of users that has to handle (or at least protect) a user's sensitive data, privacy should be a primary concern during each step of the PISA's development. We will collaborate with privacy experts in our research group to work on this question, implementing available techniques such as privacy-preserving computation of recommendations [4] to minimize privacy risks

IV. Existing Work

The PISA project uses 4 areas of existing knowledge:

- **Enterprise Security:** principles from enterprise security such as policy specification languages, access control systems, risk assessment strategies and the processes used by security operations centers are all important sources of inspiration for the PISA. The PISA project uses existing enterprise risk assessment literature here [5], but goes beyond existing strategies by adapting them for use by end-users rather than in an enterprise context. This brings with it certain challenges such as changeable environments, less regulations and undocumented processes that make securing users more difficult.
- **Existing protective measures:** existing technologies such as firewalls and virus scanners form the current set of mitigations that end-users employ to maintain their privacy and security. However, these tools have two limitations, which PISA aims to address. First, these tools and technologies focus around specific threat categories. Since PISA is to be a more general purpose security and privacy enhancing tool, it needs to be able to handle a broader scope. Second, they are usually designed to force compliance. When considering models of human reasoning such as the Reasoned Action Approach [6], it becomes clear that this is not sufficient to effect a structural behavioural change. PISA will focus on changing behaviours and attitudes instead.

- **Persuasive Technology research:** this type of research investigates how computers can be used to persuade users to a certain course of action. The PISA falls under this class of systems. It will add to this field of research, however, by studying how persuasive technology can be applied in the context of security: presently, research into persuasive technologies has focused mainly around healthcare [7].
- **Machine Learning:** to address the challenge of *dynamic threats*, the PISA incorporates existing research from the field of machine learning [8] to adapt its rules based on the context in which it operates and the preferences of its users.

V. Research Methodology

The PISA project uses an iterative prototype development method during which the functionality and scale of the experiments is increased over time. We will cooperate with our project partners, KPN[2] and XS4All[3] for development of the later prototypes and to conduct a large scale field test of our final prototype. One iteration of the PISA tool has already been implemented previously [9], which has resulted in the architecture described below. The next prototype will include functionality to measure computer ability of a user, to be validated in a lab and a field experiment that compares the results with existing computer ability questionnaires. This base functionality can then be used in the following prototypes. Throughout this process, we consult with other companies (including a bank and several consultancy firms) for expert opinion, potential case studies and feedback on our designs. The research questions posed previously can be answered through a series of experiments and consultation of literature:

- *Q1.1, adapting enterprise risk assessment methodologies for an end-user context:* This question reviews existing literature on risk assessment methodologies [5]. Using this, we can discover commonalities between enterprise and end-user contexts, which in turn can be used to devise a lightweight risk assessment methodology for end-users. This methodology can be validated by a case study using a new iteration of the PISA prototype that incorporates the risk assessment methodology.
- *Q1.2, measuring user ability and motivation:* The *ability* dimension of the question can be answered by applying the automated computer ability measurement method described above. The PISA's ability to motivate its users can be maximized by applying persuasive technology techniques described in literature such as the PSD model [10] and the Elaboration Likelihood Model [3]. The effectiveness of these techniques can then be validated using a field test on each new iteration of the PISA tool that rates the user experience of the tool at regular intervals. These tests can be used to ascertain the effectiveness at changing behaviors, perceived usability and utility, and user skill and motivation regarding security and privacy. The first prototypes will use colleague researchers as subjects, and later iterations will use field tests with a project partner, which is a telecommunications provider in the Netherlands.
- *Q1.3, visualising privacy requirements and status:* Data of the user needs to be visualised in a manner that does not overwhelm the user. As such, a tiered reporting system will be devised that aggregates individual sensor information into statistics that are usable and informative to non-expert end-users. This visualisation method will incorporate existing techniques, borrowing the visualisation results from our companion project TREsPASS[4] and testing these in our field experiments. These tests can be performed by applying an iteration of the PISA that incorporates such a visualisation method in its system in an experiment that tests usability and user experience of the PISA prototype.
- *Q1.4 structuring persuasive human-computer dialogue:* This question ties in heavily with Q2.1 and Q2.2, and as such the answer of this question can be found by applying the same method as proposed above.
- *Q2.1, identifying end-user privacy requirements:* This question can be answered by identifying threats and mitigations to an end user using risk assessment on scenarios with personas using expert opinion. These threats and mitigations can be used to elicit possible requirements. This list can serve as a basis for a *Risk Repository* which can be updated with new threats and privacy requirements as needed by security experts and the end-users of the PISA. Currently, we are applying risk assessment techniques to specific scenarios to elicit these requirements (one involving the risks of home banking, the other about risks that teleworker is exposed to).
- *Q2.2, identifying user-characterizing factors:* An initial set of factors is provided by the researcher and expanded & refined by consulting researchers that work with persuasive healthcare systems. This set is refined based on feedback of lab and field tests after every prototype, using interviews and identifying common factors in user opinion regarding PISA's interaction with the user.
- *Q3.1, ensuring privacy while analyzing user data:* Maintaining the privacy of a user while allowing for data analysis in a central location will require consultation of existing literature on privacy-preserving techniques. Here, we will use results obtained in earlier research in our group [4]. In each iteration of the PISA, its privacy and security characteristics will need to be assessed in order to preserve the privacy of the users (test subjects).

Because there are many research challenges that associated with the PISA project, we will answer them in a specific order. First, we have perfomed a literature study on persuasive technologies [11] (Question 1.4) and studied the structure of a risk respository (Question 2.1). We have then built a

[2] www.kpn.com
[3] www.xs4all.nl
[4] www.tresspass-project.eu

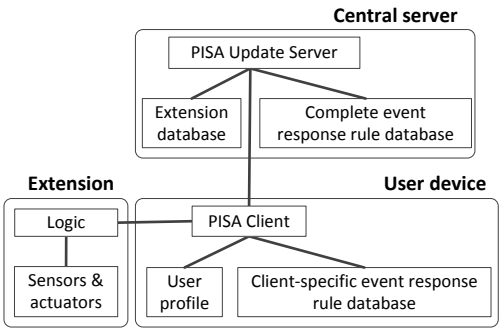

Fig. 1. The architecture of the Personal Information Security Assistant

prototype to lay the foundation for the next prototype, which will be used to identify user characterising factors (Question 2.2, Question 1.2). After this test, we will evaluate its privacy and security characteristics (Question 3.1). Question 1.4 and 1.1 are ongoing work throughout the development process and Question 3.1 will be postponed until a sufficiently advanced PISA prototype exists.

VI. Contributions

A. General

The research in the PISA project offers contributions in multiple fields:

- **Adaptive security:** it applies machine learning principles to privacy and security to deliver an evolving and personalised set of protective measures for end-users.
- **Persuading users:** the PISA project aims to improve the use of persuasive technology techniques in the field of security and privacy to investigate ways in which sustained behavioral change can be realised.
- **User characterisation:** It develops novel methods of user profiling and requirements elicitation with regards to privacy and security from the perspective of end-users.
- **An extensible software framework:** one of the project results is a software framework designed to act as a privacy enhancing technology, contributing to the social aspect of security for end-users by educating and motivating them with regards to privacy and security.

B. The PISA Software Framework

The contributions listed above are research goals of the project. As a first step to meeting these goals, we have performed literature studies, and developed a software framework for a prototype by which we can test out ideas emprically. This framework has been defined based on a refined version of a previously designed and validated PISA architecture [9]. At the core of this architecture is the concept of *extensions* that can be defined for the PISA framework, which expand the scope of protection the PISA offers. Security Experts can develop extensions that can communicate and cooperate using the PISA framework, while at the same time end-users are presented with a unified view of their security and privacy status and the mitigations available to them. An illustration of this architecture can be seen in Figure 1. It contains the following elements;

- **The PISA Client:** A program running on a user's device. Interacts with the user and uses information gathered from sensors in PISA extensions to keep a user profile up to date. Based on this user profile and a database of rules, PISA protects the user when an event takes place. It does this through advice to the user and by using actuators in PISA extensions.
- **PISA Extensions:** Plugins that can integrate with PISA, protecting the user based on a set of event-response rules associated with the PISA Extension. Sensors are programs that can monitor aspects of the user's system (such as browsing activity or typing speed) while Actuators are programs that can take specific actions within a user's system (such as starting a virus scanner). The Logic component is a program that communicates between different parts of the extension and the PISA client.
- **The PISA Update Server:** A centralized database of plugins and event-response rules that the PISA Client can use to update itself.

C. Adaptive Rules

As can be seen from Figure 1, the PISA uses event-response rules to govern when and how, it takes action. This can be to inform the user by offering advice, but can also involve independent actions such as disabling a password field in a webpage that is marked as a phishing site. The rules that accompany the extensions to the PISA software framework are created and maintained by security experts. The event-response rules consist of three elements:

- **Triggering Events:** The PISA can register events happening in a system based on the sensors that it has access to. Such a sensor could be a browser plugin that keeps track of a user's browsing activity, or a program that monitors the active processes on a device. Assuming that these sensors can register events that indicate possible security incidents in a user's environment, this possibility can be represented in terms of probability of the incident happening given the evidence (events).
- **Modifying Attributes:** The PISA uses a set of attributes that define user characteristics and preferences which modify the type of response that is triggered. An example would be a user designated as an expert receiving more technical advice than a novice user.
- **Responses:** Actions that PISA takes to inform users or protect their system's integrity. These actions are performed by calling actuators in connected PISA extensions.

By updating the *modifying attributes* of a user based on the user's activity and stated preferences, rules can automatically be adapted to fit individual users and their requirements. These modifications to the rules can then be presented to the user as a visualisation of the user's demonstrated risk appetite, to

improve the alignment between a user's risk perception and reality.

Additionally, rule feedback can be sent back to a centralised location in order to update the general template of the rule, to adapt to changing circumstances in security and privacy. For example, if a certain type of phishing becomes more popular, user feedback and registered incidents can lead to PISA assigning a higher probability of such an incident occurring by modifying the rule in the central database.

VII. Originality

The PISA project contains ideas and approaches that are original in the following areas:

- **A multidisciplinary approach to security:** the PISA proejct brings together a wide range of fields to solve a problem that still has no satisfactory answer: how do we strengthen the human link in the security chain? This approach incorporates results from other fields of research such as the social sciences to address a problem where the conventional solutions involve either purely technical or purely social mitigations. A hybrid approach may prove to be a fruitful research direction to solve this problem.
- **Enterprise security in a single-user context:** While there are specific processes, tools and experts available to a company, they are usually not applicable to a single user in a private context. The adaptation and application of enterprise security processes and tools to an end-user context has the potential to advance research into end-user privacy and security by considering existing research from a new perspective.
- **Persuasion in security:** Finally, research into persuasive technologies has so far been relatively focused on healthcare. The application of persuasive technology techniques in the context of end-user security and privacy systems may yield new insights in both persuasive technologies ánd security.

VIII. Progress

Currently, we have completed several steps that are necessary to answer the research questions that were posed in section II.

- A first prototype has been created as a proof-of-concept, yielding the architecture described above. This prototype has been used to design and validate concepts that lie at the foundation of the PISA's design. This prototype is caple of disabling a password field in a site it deems likely to be a phishing site (see prior work [9] for details). This prototype will form the basis for future PISA prototypes used in experiments.
- We have consulted previous research in persuasive technologies, risk assessment methods and policy languages to gain a better understanding of the fields in which the PISA project is to make a contribution. While background reading on risk assessment methods is still underway, sufficient literature in this area has been covered by the TREsPASS project[5]. For persuasive technologies, we have previously performed a systematic review of the available literature [11].
- The design for the PISA architecture and its event-response rules presented in brief above has been established in greater detail than described here. This serves as guide for the implementation of the coming PISA prototype.

IX. Accepted Publications

As mentioned previously, a literature review [11] and prototype implementation & Validation [9] have been published as internal reports. In addition, a paper concerning the application of behavior change support systems to the area of security and privacy has been accepted for publication in the 3rd International Workshop on Behavior Change Support Systems[7].

Acknowledgements

The PISA project is sponsored by NWO and KPN under contract 628.001.001.

References

[1] B. Bulgurcu, H. Cavusoglu, and I. Benbasat, "Information security policy compliance: An empiricial study of rationality-based beliefs and information security awareness," *MIS Quarterly*, vol. 34, no. 3, pp. 523 – A7, 2010.

[2] J. J. Gonzalez and A. Sawicka, "A framework for human factors in information security," in *WSEAS International Conference on Information Security, Rio de Janeiro*, 2002, pp. 448–187.

[3] R. Petty and J. Cacioppo, *The Elaboration Likelihood Model of Persuasion*, ser. Advances in Experimental Social Psychology. Elsevier, 1986, vol. 19, pp. 123–205.

[4] A. Jeckmans, A. Peter, and P. Hartel, "Efficient privacy-enhanced familiarity-based recommender system," in *Computer Security ESORICS 2013*, ser. Lecture Notes in Computer Science, J. Crampton, S. Jajodia, and K. Mayes, Eds. Springer Berlin Heidelberg, 2013, vol. 8134, pp. 400–417.

[5] D. Ionita, P. H. Hartel, W. Pieters, and R. J. Wieringa, "Current established risk assessment methodologies and tools," http://eprints.eemcs.utwente.nl/24541/, Centre for Telematics and Information Technology, University of Twente, Enschede, Technical Report TR-CTIT-14-04, September 2013.

[6] M. Fishbein and I. Ajzen, *Predicting and Changing Behavior: The Reasoned Action Approach*, 1st ed. Psychology Press, Jul. 2009. [Online]. Available: http://www.worldcat.org/isbn/0805859241

[7] R. Kegel and R. Wieringa, "Behavior change support systems for privacy and security," *3rd International Workshop on Behavior Change Support Systems*, to appear 2015.

[8] T. Hastie, R. Tibshirani, J. Friedman, T. Hastie, J. Friedman, and R. Tibshirani, *The elements of statistical learning*. Springer, 2009, vol. 2, no. 1.

[9] R. Kegel, "Development and validation of a personal information security assistant architecture," August 2014. [Online]. Available: http://essay.utwente.nl/65685/

[10] H. Oinas-Kukkonen and M. Harjumaa, "Persuasive systems design: Key issues, process model, and system features," *Communications of the Association for Information Systems*, vol. 24, no. 1, pp. 485–500, 2009.

[11] R. H. P. Kegel and R. J. Wieringa, "Persuasive technologies: a systematic literature review and application to pisa," http://eprints.eemcs.utwente.nl/24727/, Centre for Telematics and Information Technology, University of Twente, Enschede, Technical Report TR-CTIT-14-07, May 2014.

[5] www.tresspass-project.eu