In Search of Perfect Users: Towards Understanding the Usability of Converged Multi-Level Secure User Interfaces

Abdullah Issa
University of Melbourne, Australia
aissa@student.unimelb.edu.au

Toby Murray
University of Melbourne, Australia
toby.murray@unimelb.edu.au

Gidon Ernst
University of Melbourne, Australia
gidon.ernst@unimelb.edu.au

ABSTRACT

Converged Multi-Level Secure systems allow users to interact with and freely move between applications and data of varying sensitivity on a single user interface. They promise unprecedented usability and security, especially in security-critical environments like Defence. Yet these promises rely on hard assumptions about secure user behaviour. We present initial work to test the validity of these assumptions in the absence of deception by an adversary. We conducted a user study with 21 participants on the Cross Domain Desktop Compositor. Chief amongst our findings is that the vast majority of participants (19 of 21) behave securely, even when doing so requires more effort than to behave insecurely. Our findings suggest that there is large scope for further research on converged Multi-Level Secure systems, and highlight the value of user studies to complement formal security analyses of critical systems.

CCS CONCEPTS

• **Security and privacy → Usability in security and privacy**; • **Human-centered computing → Empirical studies in HCI**;

KEYWORDS

Multi-level security; security indicators; usable security

ACM Reference Format:
Abdullah Issa, Toby Murray, and Gidon Ernst. 2018. In Search of Perfect Users: Towards Understanding the Usability of Converged Multi-Level Secure User Interfaces. In *Proceedings of the 30th Australian Computer-Human Interaction Conference (OzCHI '18), December 4–7, 2018, Melbourne, VIC, Australia.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3292147.3292231

1 MOTIVATION & BACKGROUND

Converged Multi-Level Secure (MLS) interfaces allow users to interact simultaneously with applications that process data of different kinds, while ensuring that each kind of data is kept isolated from the others. In a Defence office context, for instance, such systems allow a user to view both the public (unclassified) Internet while writing a secret (classified) email, while ensuring that such secret data is never exposed to the Internet. Figure 1 presents a schematic of the on-screen interface of a recent converged MLS system, the Cross Domain Desktop Compositor (CDDC) [9], in this scenario.

OzCHI '18, December 4–7, 2018, Melbourne, VIC, Australia

ACM ISBN 978-1-4503-6188-0/18/12.
https://doi.org/10.1145/3292147.3292231

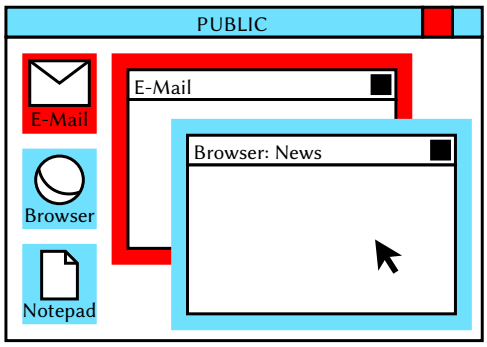

Figure 1: On-screen interface (schematic) of the Cross Domain Desktop Compositor configured with two security *domains*: PUBLIC and PRIVATE. The user is viewing a PUBLIC Internet site (light blue border), and a PRIVATE email (red border). Colours on desktop icons indicate the domain to which each belongs. The PUBLIC application has the input focus, as indicated by the top banner.

Converged MLS systems walk the fine line between, on the one hand, the need to keep different kinds of information physically separate while, on the other hand, allowing users to view and interact with each kind—and to freely move between them—on a single user interface. While their history can be traced back decades [5, 13], converged MLS systems have recently seen a resurgence in interest (see e.g. [2, 9, 15, 19, 24, 26, 28]; also modern web browsers [31, 32] and mobile phone interfaces [25]) in line with increasing security threats and the desire for greater usability.

Desktop converged MLS systems provide a converged MLS interface on a single desktop screen, keyboard and mouse, and include the CDDC [9], Qubes OS [26], Nitpicker [15], AFRL's SecureView [2] and Raytheon's Trusted Thin Client [24], amongst others. These systems present a particular challenge, since they intentionally depart from the conventions of traditional desktop operating systems, in which all user applications that appear on the single desktop screen share the same level of access to the same data. In contrast, in converged MLS systems it is common for two applications that appear on screen together to be forbidden from sharing data. Referring to Figure 1, the web browser instance (light blue border) showing the public Internet site should never obtain access to the contents of the user's secret emails (red border).

To enforce this inability of one application to access the data of another, converged MLS systems assign each running application to a *security domain*: any two applications assigned to different security domains are isolated from each other and so cannot share data. Security domains are each identified by a corresponding colour. In

Figure 1, there are two security domains: PUBLIC (whose assigned colour is light blue) and PRIVATE (whose assigned colour is red).

Because of this isolation, these systems promise very high security [9, 22]. However, it also places particular reliance upon the user. Specifically, users are required to remain aware of which security domain they are interacting with in relation to the kind of data they are inputting via the keyboard. For example, in Figure 1 the isolation between the PUBLIC web browser and the PRIVATE email client would be for naught if the user inadvertently entered secret email text into the browser's address bar.

Converged MLS systems are therefore vulnerable to mode confusion. To combat this risk, while trying to preserve the traditional desktop user interface, modern converged MLS systems implement a common set of on-screen *security indicators*. Firstly, they decorate the borders of each application window with the colour of the application's security domain, as in Figure 1. Secondly, they provide a graphical indication of the security domain of the application that currently has the input focus, which we call the *current domain*. The CDDC for instance draws a unobscurable, coloured banner at the top of the screen for this purpose, similarly to [9, 15, 16].

Converged MLS systems rest on the following user assumption.

A: At all times, the sensitivity of data entered by the user via the keyboard agrees with the current security domain.

This assumption has been *formalised* in parallel work, in which the CDDC's design and software components have been mathematically proved to enforce isolation between security domains [9, 22]. However, these proofs rely on assumption **A** being true in reality, which in turn rests on the user's understanding of the on-screen security indicators and what constitutes insecure behaviour.

Understanding how realistic these expectations are is vital to assessing the security of converged MLS systems. The effectiveness and design of security indicators for inducing secure behaviour has received much attention [1, 7, 10, 17, 18], particularly in the context of web security and phishing attacks [3, 6, 11, 12, 14, 27, 30, 33, 34]. Filyanov et al. [16] also studied the effectiveness of security indicators very similar to those of the CDDC. As with prior work on phishing, their work considers *typical* users in the context of an *adversary* who is actively trying to deceive the user into revealing sensitive information.

However, users of converged MLS systems are *atypical*, being Defence and Intelligence personnel holding security clearances, who are trained and habituated to the need to protect classified information, and aware of the consequences of, and range of penalties for, its exposure [23]. Malicious insiders notwithstanding [21], users of converged MLS systems are assumed to have a high degree of both intrinsic and extrinsic motivation to behave securely.

Further, given the increasing prominence of human error as a cause for data breaches over malicious attacks [8, 20], we argue that it is important to understand the factors affecting secure user behaviour first in the *absence* of active adversaries and deception.

In this paper we present ongoing work to better understand user behaviour in settings in which (1) users are motivated to protect sensitive information while (2) in the absence of adversaries trying to deceive them. We report on the design and preliminary results from a lab study carried out using the CDDC. Our specific research questions were:

RQ1: Do users behave securely when using the CDDC?
RQ2: Do users understand what constitutes insecure behaviour?
RQ3: Do users understand the on-screen security indicators?
RQ4: When users behave insecurely, what is the cause?

Aggregate, anonymous results are available in full at [29].

2 STUDY & EVALUATION

The study was observational: 21 participants were provided with formal training on the CDDC, before being observed while carrying out a series of tasks using the device to measure to what degree they behaved securely. Our study was IRB-approved, and participants provided written informed consent.

Experiment Scenario. The study was conducted with university students in a lab setting. This was in stark contrast to the target user population and deployment environment for converged MLS devices like the CDDC: security-conscious Defence and Intelligence personnel interacting with highly classified data.

To address the obvious mismatch we therefore devised an experiment scenario to (1) present familiar applications to the study population, while (2) increasing the likelihood that participants would be motivated to protect sensitive data during the experiment. Our chosen scenario was the Facebook social media site, under the assumption that it would be familiar to study participants due to its popularity with university students [4], while providing an environment in which participants would naturally value protecting sensitive data. (The pre-experiment questionnaire—see below—was designed to test these assumptions.)

Specifically, the CDDC was configured to operate two domains which we name here PUBLIC (whose assigned colour was light blue) and PRIVATE (whose assigned colour was red). These colour assignments followed the default configuration for the CDDC. We leave studying the influence of colour choice for future work.

To minimise risk to participants, each was given credentials to a dummy Facebook account to use for the experiment. The intent was to create a scenario in which the Facebook account and its private data represented sensitive information that should only be accessed in the PRIVATE domain, and such information should never be revealed to the PUBLIC domain. The written instructions to participants for the tasks to carry out during the observational study (see below) included the account credentials and told participants to keep the credentials secret, and that they should only ever be entered into PRIVATE applications, lest they risk being stolen. Participants were also instructed in writing here never to enter PRIVATE information into PUBLIC applications or documents.

Participants. Our study's population was 21 university students (5 female, 16 male), aged over 18, studying a range of degrees (e.g. Nursing Science, Arts, Aviation, Accounting, Psychology, Engineering, Computing, etc.), across Bachelors, Diploma and Masters level.

Pre-Experiment Questionnaire. The pre-experiment questionnaire had three primary purposes. Its first was to collect demographic information. Its second was to assess the applicability of the Facebook scenario—a key factor in our experiment's validity. To do this it asked participants closed-form questions about their usage of Facebook, to determine its familiarity to participants; yes/no questions about whether they employ common Facebook security

controls for their profile, to assess whether participants actively value the privacy of sensitive information it contains; and Likert-scale questions to assess whether they value the security of their online accounts. Thirdly, the questionnaire was designed to shed light on participants' degree of security-consciousness, for which it also asked a series of Likert-scale questions.

Training. Formal training for the CDDC was provided in the form of a short video presentation. The video presentation is available online: [29]. It was followed by an opportunity to ask questions. The intent was to educate participants not only on the CDDC but also on how to behave securely while using it. The $2\frac{1}{2}$-minute video presentation explained the CDDC's purpose, the idea of separate domains, the CDDC's security indicators and how to interpret them, how to operate the CDDC (e.g. switching between domains and applications), and the need to ensure that information does not leak from one domain to another. It gave explicit instructions on insecure behaviours to avoid, namely typing PRIVATE information into PUBLIC applications or documents.

Observational Study. For the observational study, participants were provided with a written sequence of 11 tasks to carry out in a fixed order on the CDDC. Participants took on average approximately 10 minutes to complete all tasks. These tasks were carefully designed to provide two security *decision points*, in which users would have to make implicit security decisions with the goal of observing whether they would make the correct decision at each point. The first decision point was structured so that behaving securely required less effort than behaving insecurely, while the second one was structured the opposite way. Neither involved deception, by design (see Section 1).

The experiment began with a browser running in the PRIVATE domain open to the Facebook website. On the Desktop were two icons for Notepad documents, one from each domain, as well as an icon for a PUBLIC Tic-Tac-Toe game. The coloured border around each of the icons identified to which domain it belonged.

Participants were instructed to log in to the Facebook account and perform some standard actions ("like" some pages, post a status, watch a video), in order to get habituated to the CDDC. They were then instructed to open the PUBLIC Tic-Tac-Toe game. The game gave the user the option to log-in via the Facebook account, which was the first security decision point. Choosing to log-in would open a browser window in the PUBLIC domain into which the user could enter the account credentials. Entering the account credentials here would be insecure, as it would expose the PRIVATE Facebook credentials to the PUBLIC domain. However users also had the option to skip logging-in, which required less effort. The user was instructed to save their Tic-Tac-Toe score into the PUBLIC Notepad document on the Desktop, to ensure that this document would already be open when the user reached the second decision point.

Users were then instructed to switch back to the PRIVATE Facebook window, perform another standard action (read posts in a private group) and to access some secret information, a unique 4-digit code stored in the Facebook account. Users were instructed to "Type your unique four-digit code into a Notepad document from the Desktop". This instruction purposefully avoided clarifying *which* of the two Notepad documents the user should choose. The user's choice here constituted the second security decision point. Choosing the already-open PUBLIC document—the easier choice, since the PRIVATE document would not yet have been opened—would be insecure, since the 4-digit code was PRIVATE.

Table 1: Pre-questionnaire, Likert results regarding security consciousness: 1=Strongly Disagree, 5=Strongly Agree.

Question	Mode	M	SD
I choose strong passwords	5 (38%)	4.05	0.92
I always try to behave securely online	4 (48%)	4.10	0.83
I think about my security when I am online	5 (33%)	3.71	1.15
I consider myself a security-conscious person	4 (43%)	3.57	0.93
I avoid providing my personal details online	5 (48%)	4.00	1.14

Post-Experiment Questionnaire. The post-experiment questionnaire comprised a series of yes/no and Likert-scale questions. It had two purposes: (1) to probe participants' understanding and awareness during the experiment (see Table 2), and (2) to further assess the study's validity.

3 RESULTS

User Behaviour. The primary measure of whether participants behaved securely was whether they made correct decisions at each of the decision points in the experiment. Of the 21 participants, 19 made the correct decision at both decision points (choosing not to log into the Facebook account from the PUBLIC domain and choosing the PRIVATE Notepad document to save the secret code, respectively). Of the two participants who behaved insecurely, one made both decisions incorrectly, while the other made only the second decision incorrectly.

Pre-Questionnaire. On the pre-questionnaire, 18 of the 21 participants indicated that they use their Facebook account for at least one hour per week, and 14 indicated usage $\geq$ 3 hours/week. Overall, participants assigned moderate importance to their Facebook accounts (Likert question "How important is your Facebook account to you?", where 1=Unimportant, 5=Important: mode: 3 (38% of participants), M=3.19, SD=1.03) but high importance to their online accounts in general (Likert scale question "The security of my online accounts is very important to me", where 1=Strongly Disagree, 5=Strongly Agree: mode: 5 (71% of participants), M=4.57, SD=0.81); all but two of the participants indicated they had used the Privacy Settings on their Facebook account; 14 (66%) also employed the Limited Profile Facebook feature; while 20 participants had changed the privacy settings on at least one social media account.

Results from the pre-questionnaire about general security conciousness are shown in Table 1. The participants choose strong passwords (M=4.05, SD=0.92) and avoid providing personal details online (M=4.00, SD=1.14), but indicated that they think about security when online to a lesser extent (M=3.71, SD=1.15).

Post-Questionnaire. Table 2 summarises results from the post-experiment questionnaire to understand potential causes of the insecure behaviours that were observed during the experiment.

The post-questionnaire also asked participants about the adequacy of the training. Participants overwhelmingly Strongly Agreed with the statement that "The instructions and explanations by the

Table 2: Results for understanding insecure behaviour, for the two (of 21) participants who behaved insecurely at either decision point (DP), plus population summary statistics. Numeric answers are Likert scale: 1=Strongly Disagree, 5=Strongly Agree.

Participant ID	016	019	Summary Statistics		
First DP (Chose not to log-in to Facebook from PUBLIC)	✘	✔			
Second DP (Chose PRIVATE Notepad document)	✘	✘	Mode	M	SD
I was aware at all times of which kind of application (PUBLIC vs PRIVATE) I was using	5	5	5 (62%)	4.43	0.98
While making decisions, I considered whether they would compromise secret information	3	4	4 (52%)	3.90	1.09
I believe that I made the correct decisions	3	4	5 (48%)	4.33	0.73
I felt that I was asked to enter my secret credentials where I should not have	5	1	1 (29%)	2.86	1.53
I was aware that the Facebook page held PRIVATE information not to be revealed to PUBLIC applications	3	4	5 (67%)	4.48	0.87
I believe that I chose the correct Notepad document	Yes	Yes	Yes: 95%		
I was aware that the PRIVATE Notepad document was the correct one	Not sure	Yes	Yes: 90%		

experimenters alone provided me with enough information to understand the differences between the PRIVATE and PUBLIC applications" (mode: 5 (81% of participants), M=4.81, SD=0.4).

Answers for "If I had used my own Facebook account for this experiment, I think I would have behaved differently" (Likert scale, 1=Strongly Disagree, 5=Strongly Agree) were fairly evenly spread out. The most popular answer was 1=Strongly Disagree (29% of participants), next to 4=Agree (23% of participants), M=2.81, SD=1.47.

4 DISCUSSION

The main finding is while almost all users behaved securely, we did observe two who behaved insecurely (**RQ1**). Given the small population size of 21 this fault rate is statistically imprecise, but would probably be too high in security-critical environments.

For the 19 participants who behaved securely, naturally there is little evidence to support the hypotheses that they did not understand what constitutes secure behaviour (**RQ2**), or did not understand the on-screen security indicators (**RQ3**). Indeed, given that the second decision point was designed so that the most natural course of action was to behave *insecurely*, we hypothesise that these 19 participants understood not only the security indicators but also what would constitute (in)secure behaviour.

Regarding potential causes of insecure behaviour (**RQ4**), we examine Table 2. For participant 016, when reaching the first decision point, they explicitly asked the student researcher if they agreed that the PUBLIC log-in window was "the correct one, right?", who incorrectly replied that it was. Indeed, the data in Table 2 seems to confirm that this participant was aware they might have been making the incorrect choice to log-in at the first decision point. For the second decision point, it appears as if participant 016 failed to understand that choosing the PUBLIC Notepad document was incorrect. For participant 019, based on the responses we hypothesise that inattentiveness to the CDDC's security indicators was a factor but that failing to understand how to behave securely was not.

Validity. Regarding internal validity, the results for participant 016 who behaved insecurely are possibly biased, as mentioned above. No other instances of such potential bias were present in the video observation of each participant. While we believe that no participants had trouble distinguishing the colours of the CDDC's security indicators, we did not explicitly control for this.

Our use of university students as the study population has implications for external validity. Our experiment was designed to control for this bias (see Section 2), under the assumption that Facebook would be a familiar platform to the participants and that they would value the security of their Facebook accounts. Participants overwhelmingly employ security features to limit who can view information in their Facebook and social media accounts. We conclude that participants actively value the security of private information stored in those accounts. Participants also appear to be relatively security conscious (Table 1), in line with our expectations about the CDDC's target user population. Finally, we note that the use of a dummy Facebook account might have biased our results.

5 OUTLOOK

We seek to understand the validity of the basic user assumption that underpins the security of modern converged MLS systems, focusing on the CDDC. In contrast to prior work, we considered an environment in which (1) users would be motivated to protect sensitive information, in line with typical users of such systems: Defence and Intelligence personnel; while (2) in the absence of an adversary trying to deceive users, given that human error is a major cause of data breaches over malicious attacks.

Our initial results are encouraging: all but two of our 21 participants behaved securely throughout the experiment, even when doing so required more effort than to behave insecurely. This contrasts to prior work [16], likely because of our focus on motivating secure behaviour and the absence of deception attacks.

Our findings are necessarily constrained by the study's limited scope. Likewise, our understanding of subjects' behaviour is hampered by the limited amount of qualitative information we collected.

Immediate future work will collect qualitative information (e.g. via a think-aloud protocol), while having users operate the CDDC for longer periods carrying out more complicated tasks. However even our initial results indicate that there is room to improve the design of converged MLS systems. Further, that user studies are *essential* to complement traditional formal security analyses, for properly evaluating the security of such critical systems.

ACKNOWLEDGEMENTS

Mark Beaumont, Tilman Dingler, Paul van Oorschot, Niels Wouters, and anonymous reviewers provided valuable feedback on this paper.

REFERENCES

[1] Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, Yang Wang, and Shomir Wilson. 2017. Nudges for Privacy and Security: Understanding and Assisting Users' Choices Online. *Comput. Surveys* 50, 3 (Aug. 2017), 44:1–44:41.

[2] Air Force Research Laboratory AFRL/RIEB. 2013. SecureView. https://www.ainfosec.com/innovative-products/secureview/.

[3] Devdatta Akhawe and Adrienne Porter Felt. 2013. Alice in Warningland: A Large-Scale Field Study of Browser Security Warning Effectiveness.. In *USENIX Security Symposium*, Vol. 13.

[4] Saleem Alhabash and Mengyan Ma. 2017. A tale of four platforms: Motivations and uses of Facebook, Twitter, Instagram, and Snapchat among college students? *Social Media + Society* 3, 1 (2017). https://doi.org/10.1177/2056305117691544

[5] Stanley R Ames Jr. 1977. *User Interface Multilevel Security Issues in a Transaction-Oriented Data Base Management System.* Technical Report. Mitre Corp.

[6] Chaitrali Amrutkar, Patrick Traynor, and Paul C van Oorschot. 2015. An empirical evaluation of security indicators in mobile web browsers. *IEEE Transactions on Mobile Computing* 14, 5 (2015), 889–903.

[7] Bonnie Brinton Anderson, C Brock Kirwan, Jeffrey L Jenkins, David Eargle, Seth Howard, and Anthony Vance. 2015. How polymorphic warnings reduce habituation in the brain: Insights from an fMRI study. In *SIGCHI Conference on Human Factors in Computing Systems*. 2883–2892.

[8] Ramakrishna Ayyagari. 2012. An exploratory analysis of data breaches from 2005-2011: Trends and insights. *Journal of Information Privacy and Security* 8, 2 (2012), 33–56.

[9] Mark Beaumont, Jim McCarthy, and Toby Murray. 2016. The Cross Domain Desktop Compositor: using hardware-based video compositing for a multi-level secure user interface. In *Annual Computer Security Applications Conference (ACSAC)*. ACM, 533–545.

[10] Cristian Bravo-Lillo, Saranga Komanduri, Lorrie Faith Cranor, Robert W Reeder, Manya Sleeper, Julie Downs, and Stuart Schechter. 2013. Your attention please: designing security-decision UIs to make genuine risks harder to ignore. In *Symposium on Usable Privacy and Security (SOUPS)*. 6.

[11] Alexander De Luca, Bernhard Frauendienst, Max-Emanuel Maurer, Julian Seifert, Doris Hausen, Niels Kammerer, and Heinrich Hussmann. 2011. Does MoodyBoard make internet use more secure?: Evaluating an ambient security visualization tool. In *SIGCHI Conference on Human Factors in Computing Systems*. ACM, 887–890.

[12] Rachna Dhamija, J Doug Tygar, and Marti Hearst. 2006. Why phishing works. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*. ACM, 581–590.

[13] Jeremy Epstein. 2006. Fifteen years after TX: A look back at high assurance multi-level secure windowing. In *Annual Computer Security Applications Conference (ACSAC)*. IEEE, 301–320.

[14] Adrienne Porter Felt, Robert W Reeder, Alex Ainslie, Helen Harris, Max Walker, Christopher Thompson, Mustafa Embre Acer, Elisabeth Morant, and Sunny Consolvo. 2016. Rethinking Connection Security Indicators.. In *Symposium on Usable Privacy and Security (SOUPS)*. 1–14.

[15] Norman Feske and Christian Helmuth. 2005. A Nitpicker's guide to a minimal-complexity secure GUI. In *Annual Computer Security Applications Conference (ACSAC)*. IEEE, 85–94.

[16] Atanas Filyanov, Aysegül Nas, Melanie Volkamer, and Marcel Winandy. 2013. *On the Usability of Secure GUIs.* Technical Report HGI-TR-2013-002. HGI System Security Lab, Ruhr-University Bochum.

[17] Simson Garfinkel and Heather Richter Lipford. 2014. *Usable security: History, themes, and challenges.* Synthesis Lectures on Information Security, Privacy, and Trust, Vol. 5. Morgan & Claypool Publishers. 1–124 pages.

[18] Marian Harbach, Markus Hettig, Susanne Weber, and Matthew Smith. 2014. Using personal examples to improve risk communication for security & privacy decisions. In *SIGCHI Conference on Human Factors in Computing Systems*. 2647–2656.

[19] François Lesueur, Ala Rezmerita, Thomas Herault, Sylvain Peyronnet, and Sébastien Tixeuil. 2010. SAFE-OS: A secure and usable desktop operating system. In *International Conference on Risks and Security of Internet and Systems (CRiSIS)*. IEEE, 1–7.

[20] Divakaran Liginlal, Inkook Sim, and Lara Khansa. 2009. How significant is human error as a cause of privacy breaches? An empirical study and a framework for error management. *Computers & Security* 28, 3-4 (2009), 215–228.

[21] Mark Maybury, Penny Chase, Brant Cheikes, Dick Brackney, Sara Matzner, Tom Hetherington, Brad Wood, Conner Sibley, Jack Marin, and Tom Longstaff. 2005. *Analysis and detection of malicious insiders.* Technical Report. Mitre Corp.

[22] Toby Murray, Robert Sison, and Kai Engelhardt. 2018. Covern: A Logic for Compositional Verification of Information Flow Control. In *IEEE European Symposium on Security and Privacy (EuroS&P)*.

[23] Parliamentary Joint Committee on Intelligence and Security. 2008. Review of Administration and Expenditure No. 8 - Australian Intelligence Agencies. Chapter 2: Administration. https://www.aph.gov.au/parliamentary_business/committees/House_of_Representatives_Committees?url=pjcis/adminexp8/report/chapter%202.pdf.

[24] Raytheon Company. 2014. Raytheon Trusted Thin Client. https://www.raytheon.com/capabilities/rtnwcm/groups/gallery/documents/digitalasset/rtn_216411.pdf.

[25] Franziska Roesner and Tadayoshi Kohno. 2013. Securing Embedded User Interfaces: Android and Beyond.. In *USENIX Security Symposium*. 97–112.

[26] Joanna Rutkowska and Rafal Wojtczuk. 2010. *Qubes OS Architecture.* Technical Report. Invisible Things Lab.

[27] Stuart E Schechter, Rachna Dhamija, Andy Ozment, and Ian Fischer. 2007. The emperor's new security indicators. In *IEEE Symposium on Security and Privacy (S&P)*. IEEE, 51–65.

[28] Jonathan S Shapiro, John Vanderburgh, Eric Northup, and David Chizmadia. 2004. Design of the EROS trusted window system. In *USENIX Security Symposium*. USENIX Association, 12–12.

[29] Abdullah Issa, Toby Murray, and Gidon Ernst. 2018. In Search of Perfect Users—Study Results. https://s3-ap-southeast-2.amazonaws.com/perfect-users/index.html.

[30] Joshua Sunshine, Serge Egelman, Hazim Almuhimedi, Neha Atri, and Lorrie Faith Cranor. 2009. Crying Wolf: An Empirical Study of SSL Warning Effectiveness.. In *USENIX Security Symposium*. 399–416.

[31] Shuo Tang, Haohui Mai, and Samuel T King. 2010. Trust and Protection in the Illinois Browser Operating System.. In *USENIX Conference on Operating Systems Design and Implementation (OSDI)*. 17–32.

[32] The Chromium Project. 2018. Site Isolation. http://www.chromium.org/Home/chromium-security/site-isolation.

[33] Tara Whalen and Kori M Inkpen. 2005. Gathering evidence: use of visual security cues in web browsers. In *Proceedings of Graphics Interface 2005*. Canadian Human-Computer Communications Society, 137–144.

[34] Min Wu, Robert C Miller, and Simson L Garfinkel. 2006. Do security toolbars actually prevent phishing attacks?. In *SIGCHI Conference on Human Factors in Computing Systems*. ACM, 601–610.