March 2019

Next Gen Firewalls: An Investigation into the Impact of User Activity Context on Firewall Administration

Michael Finbar Prindle
Worcester Polytechnic Institute

NEXT GEN FIREWALLS:
An Investigation into the Impact of User Activity Context on Firewall Administration

a Major Qualifying Project Report

submitted to the Faculty of

Worcester Polytechnic Institute

in partial fulfillment of the requirements for the

Degree in Bachelor of Science

in

Computer Science

by

Michael Prindle

Date: March 22, 2019

Project Advisors:

Professor Craig A. Shue, Co-advisor

Professor Lane T. Harrison, Co-advisor

Abstract

This pilot study was performed to examine the effects of presenting user activity-based data alongside standard network traffic data on network administrators' firewall configuration behavior, specifically that provided by the PEACE firewall system. Utilizing a web application to simulate interaction with a firewall, behavior was compared between traffic flows containing PEACE or only non-PEACE data. Our results were suggestive of a correlation between PEACE data and increased confidence in decision making, as well as a trend towards marking previously difficult to categorize flows as legitimate and nonmalicious. Further studies with larger sample sizes are necessary to adequately establish causation.

Contents

1 Introduction

Modern network administrators work daily with firewall technology, filtering traffic based on characteristics such as source and destination port and IP address, hostname, and time. While sufficient for policing a large subset of internet traffic, a number of popular attack vectors, such as the Microsoft Office macro attack [2] are nearly impossible to categorize without additional context. Current industry leaders such as Palo Alto Networks use technology to attempt to identify the application and user behind a given network flow, providing network administrators with additional policy configuration options [6], but this still is not sufficient for certain attacks such as the aforementioned macro attack.

The PEACE systems aims to provide additional context to network administrators. Comprising a two part system, with administration software on an admin machine interacting with software installed on all host machines in an organization, PEACE provides application installation location (path), keystroke count, mouse click count, and GUI text corresponding to on-screen behavior (e.g. "new Powerpoint pane" corresponding to the opening of a new Powerpoint window). This system allows administrators to view the behavior of a user leading up to the creation of a network connection, with the goal of enabling user intention-based policies rather than simple packet attribute-based policies.

To test the effects of PEACE data on firewall management, we created a simple web app that allows for the simulation of network administrator activity with custom network traffic flows meant to correspond with actual network behavior as it would be portrayed by the

PEACE system. We conducted trials on a number of non-expert participants, and found that the inclusion of PEACE data increased the confidence that participants felt while making the decision to block or allow a given network connection, as well as the likelihood of categorizing traffic that would be considered difficult to categorize by normal firewalls as legitimate and non-malicious. The low sample size of the trial leads to results being suggestive rather than conclusive, but the conclusions drawn are promising and justify further research into the subject matter.

2 Background

This section contains background information, meant to summarize relevant portions of the field of network security, and to provide context for some of the decisions made when assembling the protocol used for this research experiment. This includes summarizing the current goals of network administrators and the challenges they face, and explores potential ways that new technology such as the PEACE system may address some of these challenges.

2.1 Goals of Network Administrators

Modern computer networks developed from continuous revisions of technology developed in the early 1960s for sharing research notes and academic files universities and government agencies, particularly DARPA [4]. Ultimately, the Internet remains exactly that: a means of sharing files and data. The modern Internet is ubiquitous; used to transfer everything from a first grader's book report to confidential medical records and trade secrets. Today's organizations are thus heavily invested in ensuring that only the files that they wish to be transferred on their network are in fact transferred on their network. Computers can be

compromised in a variety of ways; devices can be compromised when a virus is introduced into a system, or sensitive information can be made public when a phishing site is utilized to trick an employee into believing that a malicious site is in fact a legitimate destination for data.

Network Administrators can take on a variety of roles depending on organizational structure, and are ultimately tasked with protecting an organization's data through a combination of hardware, software, and additional techniques such as user education. The techniques that are available to an organization are often specific to that very organization; while it is easy to cut down on potential infection vectors by preventing users to connect storage media such as flash drives to organization devices, certain organizations, such as libraries, have goals that necessitate allowing users to transfer their own files on to the system.

Insider threats (intentional or otherwise) are a legitimate concern to a variety of organizations, but they are not the concern of this study. We are instead focused on attacks that are meant to exploit non-malicious users of reasonable technical intelligence. Attacks such as the aforementioned phishing scheme are one possible attack vector. A network administrator means to reduce the amount of traffic on a network to the minimum possible while still allowing all legitimate, necessary traffic to continue. To this end, firewalls are one of the most useful tools available.

A firewall is a tool that can be configured to selectively block or allow traffic based on specific policies configured by the network administrator. The proper implementation of a firewall can ensure that inter-organization network activity can be monitored and specific policies enforced regardless of the behavior of other users. Firewalls present a network administrator with powerful policy creation and policy enforcement tools, as well as allowing for flagging traffic that doesn't fall under specific existing policies, allowing an administrator

to create new, pertinent policies.

2.2 Usage of Network Firewalls

The first step that a network administrator conducts when configuring a firewall for the first time is the decision to implement a blacklist-based or a whitelist-based approach. In a blacklist-based approach, traffic is assumed by default to be legitimate, and is allowed through the firewall unless it is explicitly disallowed by the blacklist. In contrast, a whitelist-based approach assumes that all traffic is guilty until proven innocent. Each of these approaches has benefits and drawbacks; in a blacklist based approach, a user can reasonably expect that they are able to access all legitimate web content unless it has erroneously been blocked by an administrator, however, a presumption of innocence leads to more threats being able to bypass the firewall as administrators race to keep up with malicious parties. A whitelist-based approach can all but guarantee that malicious traffic is blocked unless a legitimate service is itself compromised, but this comes at the cost of potentially blocking legitimate, necessary web content until it is added to the whitelist.

There is no rule that only a whitelist or only a blacklist may be used on a given network; administrators may combine the two approaches in securing their organization's network. For instance, an administrator may whitelist the .gov or .edu domain, and then selectively blacklist sites within it. Similarly, in the case of a service where the origin of connections is not consistent or easily categorized, a whitelist may be insufficient to allow the everchanging connections to access the service.

Modern firewalls present a network administrator with a variety of different characteristics for each network connection on the network, including source and destination IP addresses, ports, and host names, as well as network protocol, time, and other relevant flags. For a

firewall to be usable, this data must be presented in a concise and user-readable manner so that policies may be generated and modified to secure a network. In addition, on top of raw network traffic data, firewalls should display network trends so that network connections that are inconsistent with typical network traffic, such as an international network connection launched by a host that works nearly exclusively with organizations within its country of origin, may be examined more closely in the event that a new policy need to be created or an existing policy updated. Again, an exclusively whitelist-based approach could ensure that this "suspicious" connection is blocked, but if it were in fact a legitimate connection, e.g. an employee communicating with a client traveling internationally, a whitelist blocking this traffic could interfere with the productivity of the organization.

A firewall is ultimately a tool with two purposes: firstly, it should be able to block simple, obvious attacks such as port-scanning attacks. Secondly, it should be able to identify suspicious traffic in the hopes of recognizing a threat before it can cause any serious harm. A firewall alone is insufficient to totally protect the data of an organization, as a firewall itself cannot repair an infected host in the event that an attack succeeds, but a properly configured firewall can ideally reduce the number of attacks that actually make it on to an organization's machines and minimize the intra-machine harm that an attack may cause.

2.3 Filtering by User Behavior and Intention

Existing firewalls are able to enforce policies largely based on network packet data, such as source or destination IP address or port, domain names, etc.; these qualities are sufficient for detecting and blocking the majority of malicious traffic (and similarly, allowing the majority of legitimate traffic). However, there is a significant enough number of edge cases that companies such as Palo Alto Networks, developer of the "Next Generation Firewall" [6],

have implemented methods of determining the user and the application behind a network flow. This cuts down on the number of edge cases, but is still insufficient for detecting specific common attacks. One popular method for attack, the Word macro virus [2] consists of utilizing macro programming in Microsoft Office to infect a device via a file such as a Word document or a Powerpoint presentation. When a user clicks on a hyperlink in a Word document, Word launches a network connection; thus, a filter prohibiting Microsoft Office from opening network connections will block certain legitimate behavior. One proposed method of detecting and correctly classifying these attacks is to filter by user behavior.

The majority of "automatic" updates for modern computers can be scheduled by an administrator, allowing a network admin to cut down on unknown automatic traffic, leaving behind primarily user-driven traffic. If a user can be trusted to make informed, non-malicious decisions (as was noted to be our assumption for the purposes of this study in a previous section), the work of deciding which traffic is legitimate vs malicious can in effect be offloaded onto the user. Research indicates that this kind of "user-based" policy development can both increase the ease with which policies are developed, and increase the reusability of policies (meaning that they could be used in a greater number of situations), as well as provide for improved usability and configuration monitoring, due to "higher-level abstractions of intent" vs lower-level application behavior being the driving force behind a given policy's development and implementation [8].

2.4 PEACE

Current industry leaders in firewall technology, such as Palo Alto Networks' Next Generation Firewall, attempt to provide ease of user-driven, intent-based access control via intelligently linking network traffic with a specific application and a specific user. At current, technologies

such as these are the best performing firewall security systems. Compared to a naive firewall that simply presents packet and application data, firewalls such as Palo Alto's help provide context in addition to raw network traffic data. PEACE is a newly developed firewall technology that aims to be the first of a new tier of firewall, combining simple packet data and application data, user and application identity data, and a new level of quantitative data representing the activity of the user and the corresponding device behavior. The PEACE system, currently in development for use on Windows operating systems, includes a piece of software installed on a device which tracks a variety of useful quantitative and qualitative metrics. Specifically, the PEACE system provides the specific path of the program or application that launched a network connection, the keystrokes and mouse clicks of the user within windows of 0-5 seconds, 0-15 seconds, 0-60 seconds, 0-5 minutes, and 0-15 minutes; as well as system-provided Graphical User Interface (GUI) data leading up the the initiation of a network connection. The PEACE system allows a network analyst to examine the behavior of a user, allowing for certain traffic that was previously impossible to categorize to be categorized trivially, such as the notorious Microsoft Office macro virus attack.

3 Methodology

3.1 Creating the Web Application Study

The goal of the study was to examine and understand how the presentation of the additional data provided by PEACE affected the decision making and other relevant behavior of participants. Because the experiment was a small-scale trial study without access to field experts, we aimed to avoid requiring non-expert participants to familiarize themselves with all of the workings of a conventional firewall, so the decision was made to develop a

web application that could allow participants to simulate the specific network administrator behavior that we were concerned with investigating; specifically, the decision making when presented with suspicious flows and told to block or allow them. The web application was designed to present network flows with their associated attributes to participants, including PEACE data while relevant, and allow them to behave as a network analyst. Ultimately, the decision to utilize a web application was made due to the ease of using one across multiple machines, allowing trials to be conducted on different devices without requiring additional configuration, as well as allowing for trial to be conducted remotely.

3.2 Designing the Study

As the desire of the study was to compare and contrast participant behavior depending on whether or not a participant was presented with PEACE data, it was necessary to split the study into multiple phases. A three phase study was settled upon, consisting of two different data sets; two phases would consist of using the same flows: one with the PEACE data hidden, and one with the data visible to a participant. A third phase would use a separate dataset. For the purposes of our experiment, non-PEACE data consisted of the following fields: time, source IP, destination IP, destination port, source port, destination host, protocol, flags, and path. The decision to include path as a non-PEACE data field was made due to the existence of other firewalls, such as the Palo Alto Networks Next Generation Firewall [6] that, while incapable of providing the exact path of an application, could make a reasonably accurate educated guess as to the specific application used, effectively replicating the effects of showing the installation path. Flows with PEACE data would include data only provided by the PEACE system: keystrokes, mouse clicks, and GUI text. Keystrokes and mouse clicks would be presented in intervals from 0-5 seconds, 0-15 seconds, 0-60 seconds,

0-3 minutes, and 0-5 minutes prior to a network connection being initiated, consistent with the current workings of the PEACE system.

The three-phase design choice provided us with two opportunities: to make general comparisons between behavior of participants when presented with PEACE data vs when not presented with PEACE data, and to specifically examine if a participant would switch their decision when presented with PEACE data for a flow that they had previously seen with PEACE data, e.g. switching from "block" to "allow" on a flow representing what could potentially be interpreted as a Microsoft Office macro attack once the presence of participant interaction was made clear by the PEACE system. This comparison necessitated that segment with the PEACE-version of the 7 flows that would be presented twice be presented *after* the non-PEACE version, so this segment was always presented last. In addition, we randomized the presentation of the first two segments (the non-PEACE version of the flows that would be presented twice, and the flows that would be presented once, always showing PEACE data). This was in the effort to present a systematic bias from occurring if participants were always presented with non-PEACE flows or always with PEACE flows first, as their commentary when PEACE data was *removed* was equally interesting to us as was their commentary when PEACE data was added in the first place. In addition, we hoped that in doing so, we would prevent the trials from being skewed by all participants developing a methodology of blocking or allowing flows exclusively utilizing non-PEACE data that they would then carry into the PEACE flows, e.g. using only the path and destination host even when eventually presented with keystrokes, mouse clicks, and GUI text.

3.3 Designing Flow Page Display

To design the web application, we examined industry leader Palo Alto Networks' Next Generation Firewall [6] to develop a sense of how firewalls presented the data. We needed to develop an easily-understood method of displaying all 21 flows (the number 21 was chosen after some initial testing on early implementations of the web app was implemented, but there was always an intention to have at least 15 flows, so the final number did not have a meaningful impact on initial flow page design). Initially, we attempted to present all of the flows in a single table as was industry standard, but ultimately the decision was made to display each flow independently. This decision stemmed from the fact that being able to parse multiple flows at once and draw conclusions that drew from several different flows was thought to be too difficult for our non-expert participants, as well as the fact that implementing mutli-flow scenarios could needlessly increase the amount of time taken by participants to conduct the study. In addition, by consolidating all of the data for a single flow onto one page, rather than having it split across multiple pages due to the presence of several flows on a single page, we hoped to gain more specific insights into the decisions made by participants on a flow-by-flow basis, and prevent participants from falling into a pattern of grouping multiple similar flows together and making the same decision for all of them despite potentially difficult to spot meaningful differences.

The ultimate final design consisted of a single page (sometimes small enough to be viewed without scrolling depending on the length of GUI text data for PEACE-enabled flows) presenting the non-PEACE elements and the PEACE elements. This design can be viewed in figure 1. The relative location of elements was ensured to remain consistent between flows by placing the GUI text (when relevant) at the bottom of the page, as it was the only element with a chance to significantly vary in size. In addition, the placement of the non-PEACE

Figure 1: Flow Page (with PEACE data present), consisting of non-PEACE data, keystroke and mouse click data, GUI data, and decision selector/feedback section

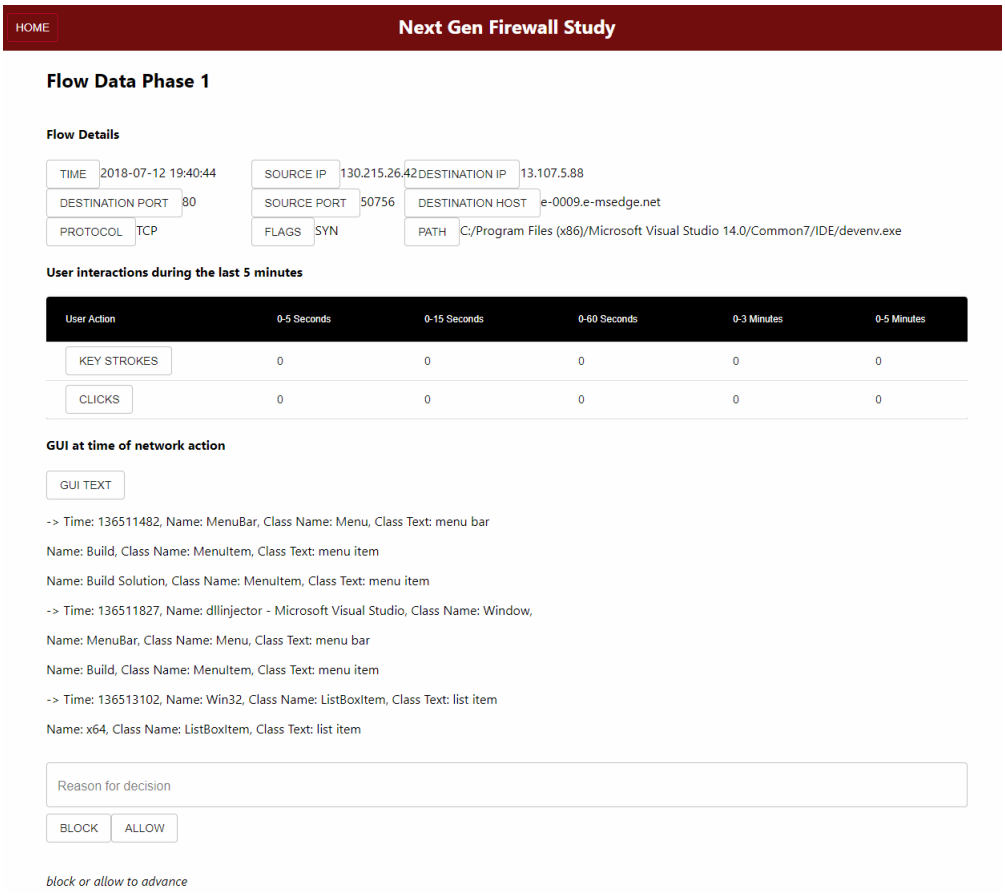

data first meant that its location would not be dependent on PEACE data being enabled or disabled.

The non-PEACE data was presented in a 3x3 grid at the top of each flow page, allowing participants to familiarize themselves with the placement of each flow component in an optimized grid format that would be unchanging across flows. Additionally, by utilizing a grid format rather than simply listing attributes without context, as is the case with some primitive firewalls, participants with less expertise were immediately able to identify each component, i.e. there was never any confusion over the source or destination IP as each was

Figure 2: non-PEACE components of flow data

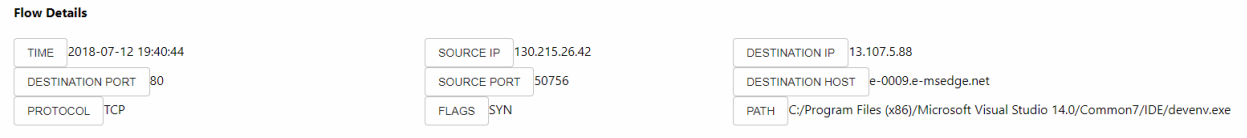

labeled.

In order to gather quantitative metrics on a participant's behavior without relying on asking questions as they proceeded through the study, each indicator label functioned as a clickable box that would highlight the indicator as being a component of a participant's decision making on a given flow. The close placement of an indicator's label next to the indicator itself also ensured that the ability to signify an indicator was always in view when a participant was utilizing a given indicator, so as the reduce the chances that they forgot to highlight the indicator. The aforementioned 3x3 grid layout, including the clickable indicator labels, can be viewed in figure 2.

Figure 3: Keystrokes and Mouse Clicks components of PEACE data

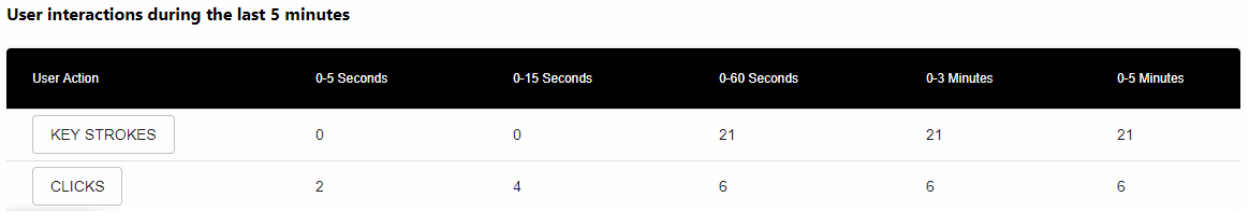

As mentioned previously, consistent with the current implementation of the PEACE system, keystrokes and mouse click data would be presented in a table with space for cumulative totals for 0-5 seconds, 0-15 seconds, 0-60 seconds, 0-3 minutes, and 0-5 minutes prior to the initiation of a network flow. To present this data, we implemented a table with keystrokes and mouse clicks sharing a single table, including a clickable indicator identical to those

utilized in the non-PEACE data segments. An example of this table can be seen in figure 3. The values in the table, as indicated by the column labels, are cumulative; for instance, in figure 3, there are a total of 21 keystrokes and 6 mouse clicks in the 5 minutes leading up to the given flow opening a network connection, with 21 of the keystrokes and 6 of the mouse clicks having been within the minute leading up to the network connection, and 2 of the mouse clicks and none of the keystrokes having been less than 5 seconds before the network connection was launched. Just as with the non-PEACE data, the goal was to keep this data all visible close to the clickable indicators, again to prevent a participant from forgetting to select the indicator.

Figure 4: GUI text associated with PEACE data. In this case, the activity presented is interaction with Microsoft Visual Studio

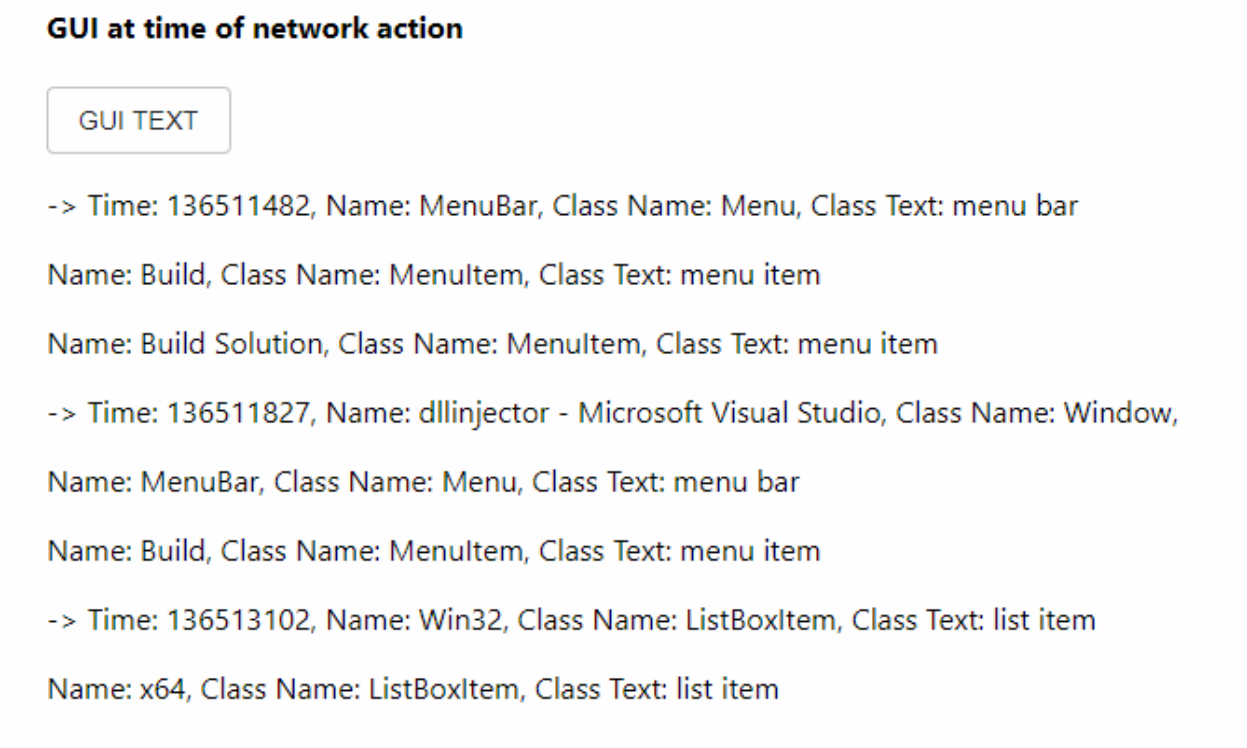

The final piece of data presented to participants was the GUI information as parsed by PEACE. This data is presented in a hierarchy consisting of fields of time, name, class name, and class text. Initially, this data was taken from PEACE as a long, barely-legible string. The decision was made to split it into a hierarchy separated by the "->" symbol before each time stamp, as discrete elements each had their own time stamp. This was perhaps the most difficult data to present in a human readable way, as the GUI text component of PEACE relies on a Windows application's self-reporting of GUI elements, which are inconsistent across applications. For instance, some applications don't utilize the "class text" field, while others don't utilize the "class name" field, and others still utilize both. Figure 4 demonstrates a single instance of GUI text; in this example, the text appears to indicate that a menu bar was interacted with, followed by a "dllinjector" window generated by Microsoft Visual Studio, followed by a "list item" of name "Win32". This is clearly difficult to parse, and even the research team continues to struggle to interpret it correctly at times, but it is an important and actively developed component of the PEACE system and was included for that reason. As with the non-PEACE data grid and the keystrokes and mouse clicks table, there is a "gui text" label indicator that may be clicked by participants.

Figure 5: Final decision selection and comment box. Participants were encouraged to explain their reasons for making a decision)

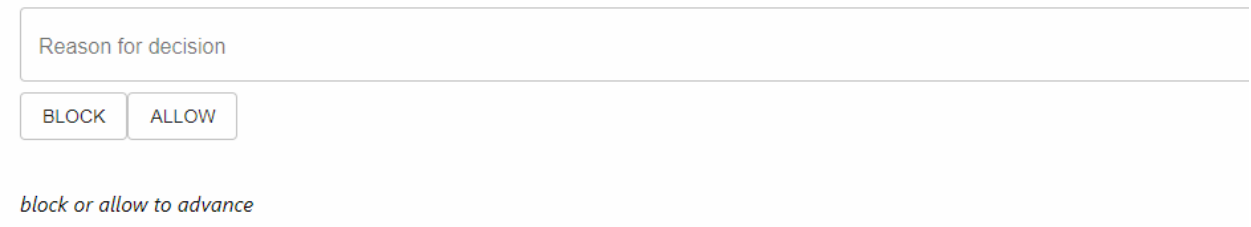

At the bottom of each flow page is a decision panel, enabling a participant to decide to

ultimately "block" or "allow" a network connection. In addition, there is a comment box that allows a participant to write down their more abstract reasoning for making a decision beyond simply the indicators used. This segment of the display page is visible in figure 5. A participant is required to ultimately make a definitive decision to block or allow network traffic, even if they are not confident about their decision, consistent with actual network administration. This prompts participants to think carefully about their decisions on difficult flows, rather than simply "skipping" a flow if a non-decision option were allowed.

3.4 Assembling Network Flows

The key component of our web application, beyond just the design, was the network flows themselves. Every network packet has certain data associated with it, such as source and destination IP address and host, and a time stamp, among other fields. These are the attributes that comprise the non-PEACE component of a flow. The PEACE system adds additional information consisting of the keystroke and mouse click counts leading up to a network connection, the GUI text, and the application path location (included in non-PEACE for for reasons stated above). These flows simulate the bulk of the data that a real network analyst would utilize to make decisions for allowing or blocking traffic, as well as creating or editing policies. The research team was provided access to a database of flows produced by the PEACE system during testing, so the obvious choice for assembling flows was to take existing flows that suited our needs and modify them when necessary before implementing them with the web application.

Certain network flows, while regularly present in an actual network, were of less interest to the research team than others. For instance, automated inbound connections such as those generated by Microsoft updates would not have any associated PEACE data, so only one

or two flows of this type were included to get a better sense of the general thought process of a participant. Flows that included PEACE data were obviously prioritized, thought this did not necessarily mean that only flows with high numbers of mouse clicks and keystrokes, and large blocks of GUI text, were included; network connections with sparse accompanying PEACE data were also important. However, due to study being a pilot study more focused on understanding the effects of PEACE data on participant thinking and decision making and less focused on drawing statistically significant qualitative conclusions, the majority of flows were selected from existing flows with higher amounts of PEACE data.

We desired to gather flows that simulated as wide a variety of network activity as possible, including both legitimate traffic and traffic that would be associated with malicious behavior; however, we also wanted to keep the flows *relatively* similar so that participants could identify key similarities and differences that would only be discernible through the use of PEACE data; for instance, the non-PEACE data for a legitimate Microsoft Office hyperlink network connection and that associated with a macro attack would be highly similar, with significant differences in PEACE data. As the flows gathered from the PEACE database were ultimately plain text that the research team was free to manipulate, we were able to modify a number of flows to simulate the behavior we desired; for instance, we were able to take a flow associated with legitimate Microsoft Word usage and modify it to reflect a macro attack.

Ultimately, this approach allowed us to encompass most, but not all types of attacks. A number of attacks, such as port-scanning attacks, would require multiple flows to represent, and multi-flow decisions were intentionally omitted from this trial. Port-scanning attacks and attacks with similar mechanisms are already blocked easily enough with existing primitive firewall tools, so the research team felt that the ability to demonstrate multi-flow traffic was not worth the additional level of complexity that it would bring to the experiment, as we

were concerned that low-expertise participants would be overwhelmed.

3.5 Building the Web Application

The web application [5] was built using the Nodejs [1] Javascript backend framework, as well as Facebook's React [7] client side framework, and deployed via Herokuapp [3] so as to be easily accessible by the web. React was used as it allowed us to run the entirety of the application on a client's device, without having to rely on managing external databases or other server side components besides deploying the app on Heroku.

When a participant opens the web app for the first time, they are presented with a brief study of the purpose and goals of the experiment, before being prompted to undergo a tutorial to familiarize them with the web application. This tutorial displays each of the sections of a flow page in the same sequential order that they were explained in a previous section, enabling a participant to understand the mechanisms of the web application. Once the tutorial is complete, a participant begins the core component of the trial where they will categorize flows as malicious or legitimate.

Each individual flow in each phase maintains its own web individual web page, and participants navigate through the application by selecting a decision to block or allow a flow, which reveals a "continue" button to the participant. In between each phase of the trial is a simple page encouraging participants to take a brief break if required. As detailed in a previous section, the presentation order of the first two sections is randomized; the labeling of "Phase 1" and "Phase 2" remains logically consistent (1 comes before 2) though if a participant were to examine the associated URL closely, they may notice that they might start on flow 8 rather than flow 1; this has no meaningful impact on the study.

Once all 21 flows have been examined, a participant would arrive at a final screen prompt-

ing them to note any final thoughts or observations that they felt would be useful to the experiment investigator. In addition, the final screen included a download button that would enable an investigator to save a file containing the qualitative metrics gathered by the web application. This file would contain time stamps for each flow, as well as the flow number, whether or not it had PEACE data exposed, whether or not it was ultimately blocked, and the indicators selected for the given flow.

3.6 Conducting the Study

Once the web app was completed and populated with flows, participants were assembled for testing. participants were not screened beyond being a frequent computer participant, as the pilot-study nature of this research experiment led to the team desiring to examine behavioral trends associated with the PEACE data presented in the flows, for which a non-expert would be sufficient for testing.

Each trial consisted of a participant being recorded completing the security simulation presented by the web application. For all but one trial, both audio and video, in the form of a screen capture, was used to gather data for later analysis. Due to technical failures in one trial, only audio was recovered. At the beginning of each trial, the participant was asked to complete the tutorial that preceded the security simulation, during which time they were free to ask the trial administrator clarifying questions. In addition, participants with less of a background in network security were provided with information that would be common knowledge to someone in the field; for instance, a list of common domains that would be familiar to experts in relevant fields, such as the name "Akamai" being that of a large content delivery company. This information was provided in response to feedback in initial trial runs, where non-experts were focusing on foreign-sounding names such as the

aforementioned "Akamai". Care was taken to emphasize that the appearance of a seemingly legitimate name did not guarantee the legitimacy of traffic. The goal of this additional information was to reduce the impact of utilizing non-experts in the study.

Once the tutorial was completed, each participant worked through the 21 flows in 3 phases; the first two phases were ordered randomly. These two phases consisted of one phase of 7 flows with PEACE data and one of 7 without; the third phase consisted of the 7 non-PEACE phases presented again, this time with PEACE data. Each participant was asked to take time to consider each field before making their decision, and to attempt to think aloud during the study, so that a combination of the flow indicators, the explanation box, and recorded audio could be compiled to attempt to understand the thinking of each participant as they progressed through any given flow.

After all 21 flows had been completed, participants were asked for any thoughts that they felt the research team should know. After allowing for open ended responses, participants were then asked for their thoughts on the PEACE vs non-PEACE phases, if they had not already provided feedback on the matter.

4 Results

4.1 Overall Indicator Trends

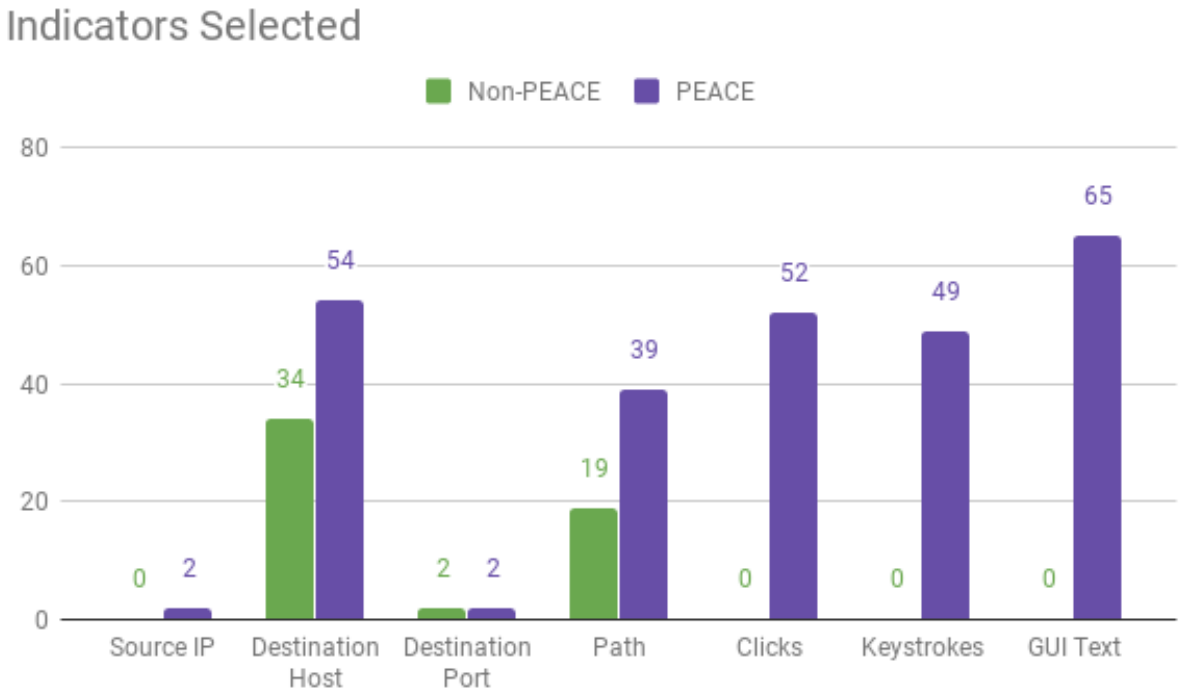

(a) Indicators selected across all flows

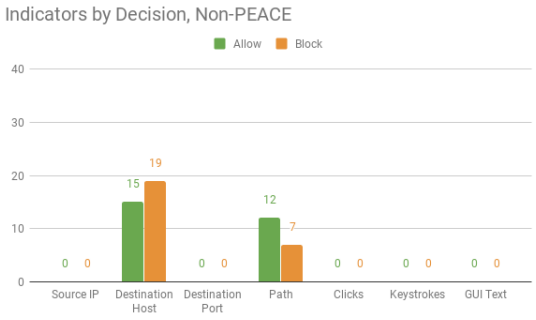

(b) Indicators selected for non-PEACE flows

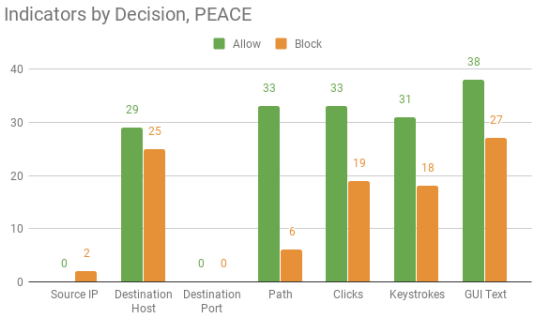

(c) Indicators selected for PEACE flows

Figure 6: Indicators across all expertise levels and participants

Per Figure 6a, across the entire participant population, destination host and path were the most popular indicators for Non-PEACE flows. In PEACE flows, these two indicators maintain popularity, coupled with similar or greater popularity of the PEACE indicators of keystrokes, clicks, and GUI text.

As seen in Figures 6b and 6c, destination host was far more popular that path when blocking a flow, both with non-PEACE and PEACE flows. When presented with PEACE data, destination host was more popular than all PEACE indicators barring GUI text in the case of a Block decision, while it was less popular than all PEACE indicators in the case of an Allow decision. Allow decisions were more frequent in flows suggesting higher user activity vs automated services. For example, an automated update by Windows Defender will have no corresponding user interaction, compared to a Google Chrome search which will include mouse clicks for opening Chrome and keystrokes for typing an address and searching it. This lack of user activity prevented a participant from quickly identifying a network flow as originating from legitimate user activity.

4.2 Participant and Expertise Trends

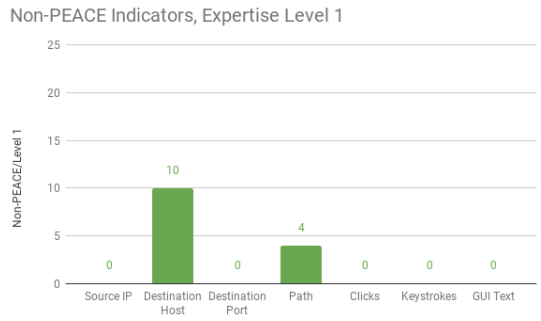

(a) Indicators selected by expertise level 1 participants, non-PEACE

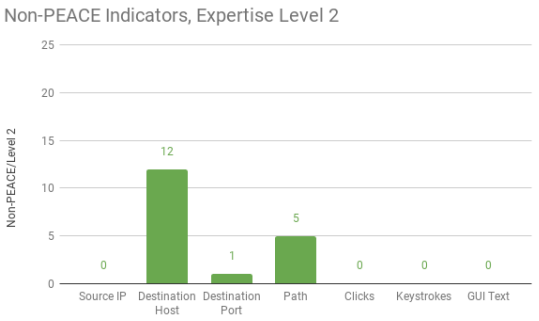

(b) Indicators selected by expertise level 2 participants, non-PEACE

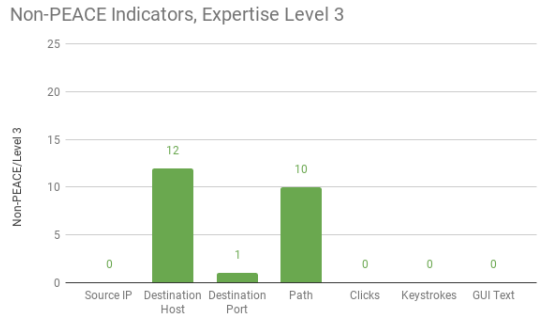

(c) Indicators selected by expertise level 3 participants, non-PEACE

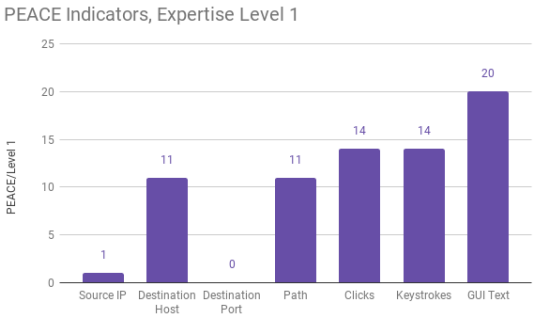

(d) Indicators selected by expertise level 1 participants, PEACE

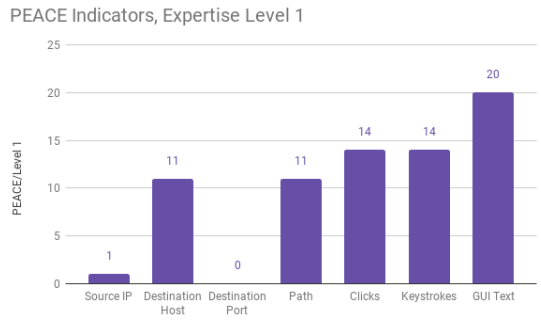

(e) Indicators selected by expertise level 2 participants, PEACE

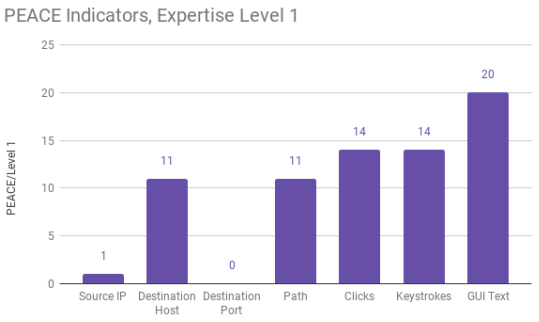

(f) Indicators selected by expertise level 3 participants, PEACE

Figure 7: Indicators separated across expertise level and PEACE vs non-PEACE

From Figure 7, based off of the hypothesis that a participant with more expertise understands what they are doing better, less experienced participants were overly reliant on destination host indicators when presented with non-PEACE flows, under utilizing the path indicator. Inversely, lower expertise participants were likely to under utilize destination host indicators when presented PEACE data based on raw counts of indicator selection, and favored GUI text indicators, whereas more experienced participants used a mix of all relevant indicators, suggesting a greater ability to synthesize the information presented by multiple indicators.

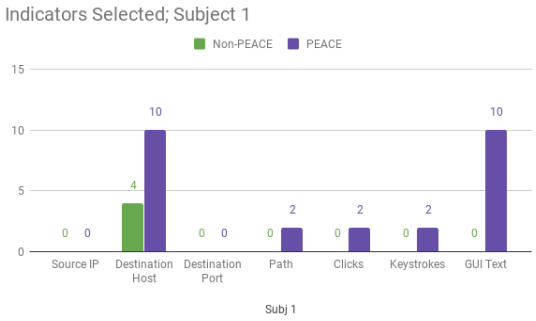

(a) Indicators selected by participant 1

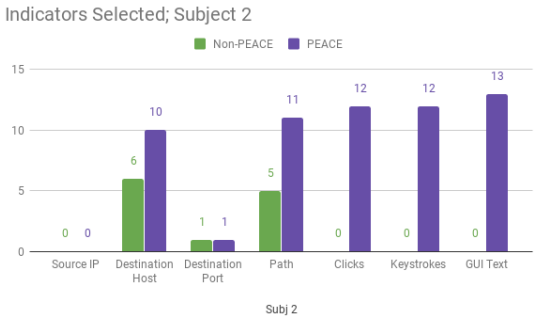

(b) Indicators selected by participant 2

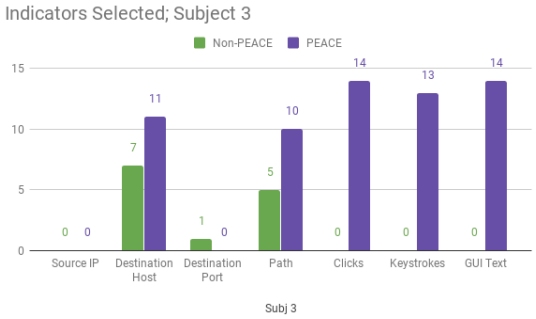

(c) Indicators selected by participant 3

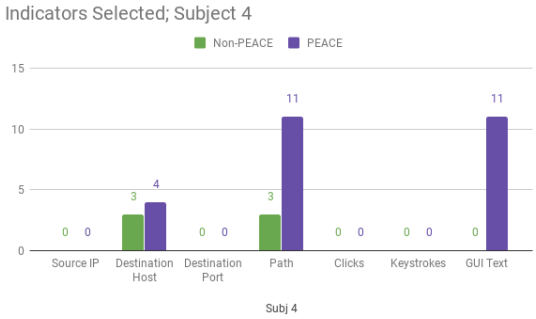

(d) Indicators selected by participant 4

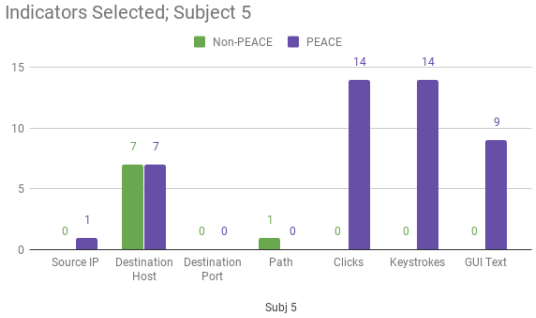

(e) Indicators selected by participant 5

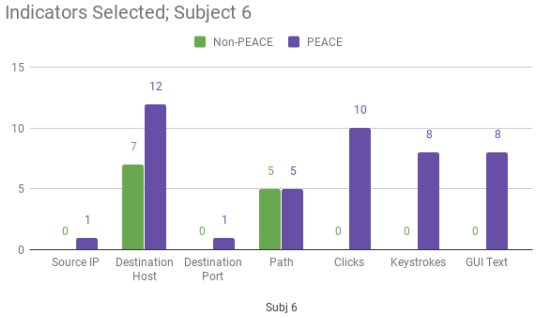

(f) Indicators selected by participant 6

Figure 8: Indicators separated across participant

From Figure 8, relative to the majority, participants 1 and 5 underutilized the path indicator (usages of 2 and 1 respectively, compared to a median of 10 selections), while participants 1 and 4 underutilized the clicks (usages of 2 and 0 respectively, compared to a

median of 10 selections) and keystrokes (usages of 2 and 0 respectively, compared to a median of 8 selections) indicators. The destination host indicator and the GUI Text indicator were both utilized universally across all participants.

4.3 Switch Indicators

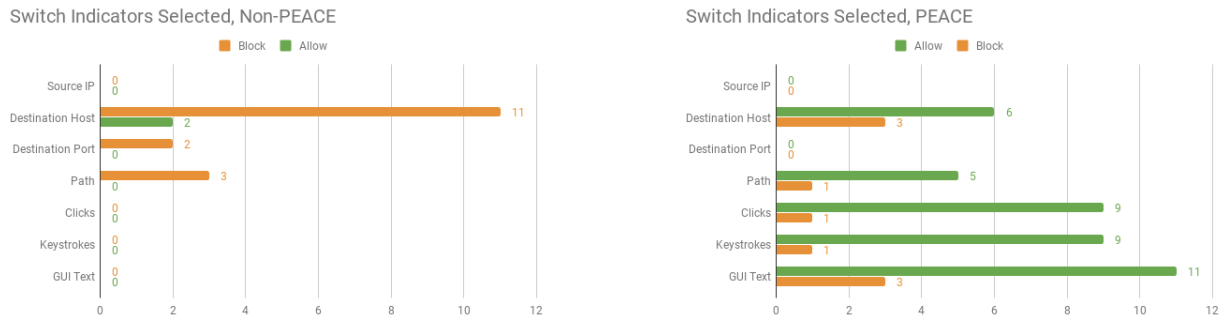

(a) pre-switch indicators, split on PEACE and non-PEACE

(b) post-switch indicators, split on PEACE and non-PEACE

Figure 9: Switch indicators split on PEACE and non-PEACE

Switch data, present in Figures 9a and 9b, demonstrate that in switch scenarios, non-PEACE block decisions heavily favored destination host indicators, though also incorporated Port and path indicators, Conversely, non-PEACE allow decisions were based exclusively on destination host. However, in block to allow switches, PEACE indicators of keystrokes, clicks, and GUI text were all utilized in greater numbers than destination host or path; in addition, path was utilized nearly as much as destination host. Switches from allow to block were based nearly exclusively on a combination destination host and GUI text data, with path, keystrokes, and clicks being barely used. Destination port was never selected as a post-switch indicator.

5 Discussion

5.1 Participant Decision Making Confidence

Some PEACE indicators (clicks, keystrokes, and GUI text) are utilized heavily when presented, along with destination host and path. Destination host and path are the primary indicators used for making decisions in non-PEACE flows, indicating that users utilized PEACE data to supplement the existing data while making decisions, rather than using it exclusively for decision making. One participant noted that "[they] felt more confident in [their] decisions once [they] had additional information that corroborated with [their] suspicions...the keystrokes and the clicking counter, along with the GUI text, was very helpful" (participant 3). Though there were slight variations in indicator selection across expertise level and participant, all data splits showed a trend of relying heavily on PEACE data when it was presented. "Having [the PEACE data] was very useful...having more data made it easier to follow, logistically, what each flow component meant" (participant 6).

5.2 Decision Switches and User Activity

PEACE data affects the decision making of the participants. Per figure 9 there were only 4 switches from allowing to blocking a trace, while there were 3 times as more, 12, switches from block to allow. One participant's observation may suggest why this is: "I thought that the GUI text being presented was very useful because it helped me have a sense of what a user was interacting with on the screen and whether it was normal or notâĂę[it] helped me understand if they were in a normal activity like looking for something to buy, based on what they were moving or clicking ...", the participant noted that "having an opportunity to see what the user was seeing and doing helped was very helpful" (participant 2). The data and

commentary of participants suggest that they would use the PEACE data to make decisions based on user activity, rather than network activity; this is exactly what the PEACE system hopes to achieve.

6 Conclusions

The study is ultimately suggestive, rather than conclusive, due to the small sample size of 6. Despite this, the preliminary results are encouraging; subjects unanimously expressed in post-trial interviews that they found PEACE data helpful for decision making, boosting confidence and ease of drawing conclusions. There are a number of potential benefits here: PEACE data may help enable existing network analysts to streamline their decision making process, incorporating user activity into their network policies. Alternatively, PEACE may lower the barrier for entry for new network analysts, decreasing the amount of technical training needed to effectively monitor a network by making network monitoring more based on familiar concepts such as user intention rather than more complex concepts. Ultimately, a larger trial, involving expert, professional network analysts is necessary to draw statistically significant conclusions, but the promising results of this study can easily justify the establishment of such a trial.

The primary contribution to relevant fields by this pilot study is the experiment designed for the study. The application and presentation of the trial allows for non-experts to simulate behavior of network administrators without needing to familiarize themselves with an actual firewall system, enabling conclusions about policy-making behavior to be drawn from a larger population at the cost of less expertise per participant. Tweaking certain values in the application, such as custom flows testing for specific subsets of attacks, may allow future

researchers to investigate other questions sharing the knowledge domain of this study without the need for developing a new custom tool.

References

[1] Node.js Foundation. *Node.js*. URL: https://nodejs.org/en/.

[2] *Frequently asked questions about Word macro viruses*. URL: https://support.microsoft.com/en-us/help/211607/frequently-asked-questions-about-word-macro-viruses.

[3] *Heroku*. URL: https://www.heroku.com/.

[4] Barry M. Leiner et al. "A Brief History of the Internet". In: *CoRR* cs.NI/9901011 (1999). URL: http://arxiv.org/abs/cs.NI/9901011.

[5] *Next Gen Firewall Study*. URL: https://next-gen-firewalls.herokuapp.com/.

[6] *Palo Alto Networks – Global Cybersecurity Leader - Palo Alto Networks*. URL: https://www.paloaltonetworks.com/.

[7] *React – A JavaScript library for building user interfaces*. URL: https://reactjs.org/.

[8] Jeffrey Shirley and David Evans. "The user is not the enemy". In: *Proceedings of the 2008 workshop on New security paradigms - NSPW 08* (2008). DOI: 10.1145/1595676.1595683.