

March 2019

Influence of End User Activity on Firewalling Decisions

Ryan William Wittenberg
Worcester Polytechnic Institute

Follow this and additional works at: <https://digitalcommons.wpi.edu/mqp-all>

Repository Citation

Wittenberg, R. W. (2019). *Influence of End User Activity on Firewalling Decisions*. Retrieved from <https://digitalcommons.wpi.edu/mqp-all/6761>

This Unrestricted is brought to you for free and open access by the Major Qualifying Projects at Digital WPI. It has been accepted for inclusion in Major Qualifying Projects (All Years) by an authorized administrator of Digital WPI. For more information, please contact digitalwpi@wpi.edu.

Influence of End User Activity on Firewalling Decisions

A MAJOR QUALIFYING PROJECT

Submitted to the Faculty of the

WORCESTER POLYTECHNIC INSTITUTE

In partial fulfillment of the requirements for the

Degree of Bachelor of Science in

Computer Science

by

Ryan Wittenberg
March 22, 2019

Professor Craig A. Shue, Project Advisor

Professor Lane T. Harrison, Project Advisor

Abstract

One prominent role of the modern-day network security analyst is firewall policy configuration and management. Despite continual advances in technology, the current extent of user activity-driven firewall management consists of associating complex network data with machines, applications, and usernames. The Policy Enforcement and Access Control for Endpoints (PEACE) system attempts to advance modern firewall technology by allowing a network analyst to associate user activity like keystrokes, mouse clicks, and graphical user interface (GUI text) with network flows, as well as providing more detailed application installation data, by installing an application on organization devices that will report relevant data to the firewall. To test the effectiveness of this new technology, this study records and compares participant behavior when presented with standard network flows verses flows with additional PEACE-exclusive insight. The initial conclusions suggest a positive impact on analyst confidence derived from the introduction of PEACE data, and justify interest in conducting a future large-scale study to more rigorously examine this area.

Contents

1	Introduction	3
2	Background	4
2.1	What is PEACE?	4
2.2	What is a firewall?	5
2.3	How do network analysts use firewalls?	5
2.4	Filtering by intention	8
3	Methodology	9
3.1	Designing the study	9
3.2	Designing the flow page display	10
3.3	Building the web application	13
3.4	Gathering network flows	16
3.5	Conducting the study	17
4	Results	19
4.1	Indicator utilization	19
4.2	Blocked flows for the non-PEACE phase by participant	20
4.3	Blocked flows after introduction of PEACE insight by participant	22
4.4	Decision changes for each flow by participant	22
4.5	Direction of decision change for each switch	22
5	Discussion	24
5.1	PEACE indicator utilization	24
5.2	Switched verdicts as a result of PEACE	24
5.3	Considerations for future studies	25
5.4	Conclusions	25

1 Introduction

Policy Enforcement and Access Control for Endpoints (PEACE) is a new firewalling system that presents a network analyst with new data metrics to help them make informed decisions when protecting a network. There are many different implementations of enterprise firewall systems, but all differ in terms of the data they provide the analyst, types of policies they can create on the network, and usability of the user interface. This study will determine which information is most helpful to a network analyst, and whether or not end user specific information has a large impact on a network analysts decision to block or allow traffic.

As technology advances, there will always be a need for improved security. In particular, as networks and malicious network attacks become more advanced, there will be an increased need for advanced firewalling techniques and traffic analytics. The PEACE system presents more information to a network analyst than ever before, but does the new information collected by the system actually help a network analyst better perform their job? This study will be the first to systematically identify which, if any, of the new information collected by the PEACE system will ultimately help advance network security by making network analytics easier.

By watching and measuring how different network analysts use different firewall systems we should be able to isolate the components of the firewall systems that are most effective and determine if the PEACE implementation is quantifiably superior to the industry standard due to its ability to dive deeper in to end user activity. For example, by isolating the data presented by both the PEACE system and industry standards and measuring how different analysts respond to the different data sets, we should be able to see how or if the different data affects network traffic classification.

2 Background

As mentioned previously, this study aims to determine the impact of the new information provided by the Policy Enforcement and Access Control for Endpoints (PEACE) system on the decision making of network analysts. 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 to attain this goal.

2.1 What is PEACE?

An understanding of PEACE technology is relevant to this study because it presents network analysts with information that has not yet been proven to be an asset when trying to protect a network. Current industry leaders in firewall technology, such as Palo Alto Networks' Next Generation Firewall[5], attempt to provide ease of user-driven, intent-based access control via intelligently linking network traffic with a specific application and a specific user. Currently technologies such as these are the best performing firewall security systems. Compared to a naive firewall that simply presents packet and application data, which we will refer to as a Tier 1 firewall, firewalls such as Palo Alto's help provide context in addition to raw network traffic data; we will refer to these as Tier 2 firewalls. PEACE is different than the current industry standard because it provides insight to help more accurately determine the intentions of the source of the network traffic by collecting data on the user's machine [4]. PEACE is a newly developed firewall technology that aims to be the first Tier 3 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 to 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. Take, for instance, the previously mentioned infected Microsoft Word attack. In the case of a legitimate connection opened by clicking on a hyperlink, PEACE would show a number of mouse clicks in the seconds before the connection was opened, as well as GUI data referencing the interaction with a hyperlink. Conversely, if the connection were an automated "phone-home" by a malicious program hidden within a document, there would be a lack of corresponding mouse clicks, and the GUI data would not line up with the expected behavior. What was previously "Bob from accounting made a connection from Word" now includes critical contextual data, allowing a network analyst

to easily distinguish legitimate traffic from malicious connections.

2.2 What is a firewall?

Understanding how firewalls work is a critical part of this study because their functionality influences the information they present to the network analysts. The first of what would be recognized today as a firewall was developed at AT&T Bell Labs by Bill Cheswick and Steve Bellovin in the late 1980s [2]. These firewalls were simply packet filters that would filter packets based on a preconfigured set of rules. From there, firewalls evolved to operate up to layer 4 of the Open Systems Interconnection Model (OSI Model) through retaining a packet until further information can be used to verify whether a packet is malicious or not; this technology was developed by Dave Presotto, Jardanan Sharma, and Kshitji Nigam, also of AT&T Bell Labs, from 1989-1990^{notes1}. The most recent significant advancement in firewall technology was the Application Layer Firewall, developed by Marcus Manum, Wei Xu, and Peter Churchyard in 1993 under the name Firewall Toolkit in 1993. Modern firewalls are simply extensions of this technology, examining the input, output, and/or access to an application or service to make a decision on the legitimacy of network traffic involving it. A firewall is ultimately a tool that captures internet traffic and examines its key characteristics in an effort to determine its purpose and whether or not it should be allowed into/out of a network. Firewalls utilize a variety of filters in order to automatically block or allow network traffic; a network security analyst is tasked with configuring those filters to the needs of their organization, and inspecting suspicious, hard to classify packets in order to help develop more accurate filters. Because firewalls generally are limited to looking at network traffic from the outside in, the information that a network analyst uses to make decisions about traffic is not able to measure the intentions of the source of the traffic. PEACE offers a new set of information to be analyzed on a network by running directly on the user's machine and collecting user actions at the time of the origin of the network traffic. This allows PEACE to gain insight on the source's intentions because of its ability to inspect network traffic from a lense within the origin machine.

2.3 How do network analysts use firewalls?

When protecting a network, the administrator needs to be able to proactively configure a firewall to prevent threats to the network, and needs to have the ability to manage threats reactively once a threat is identified. When setting up a firewall, a network admin will need to decide if the activity on the network is assumed to be harmful or safe. Assuming every process or connection is harmful by default will result in a safer environment, but it will come at a greater cost of both computer resources and the end users time. The default assumption of innocent until proven guilty, on the other hand, is quicker for the end user, but comes at the cost of less effective security. A more proactive approach to network security must be taken when unidentified data is assumed to be malicious, and

a more reactive approach must be taken when unidentified data is permitted by default.

In order to proactively manage a network, a network admin needs to have a tool to intuitively create policies regarding the data flow on the network. These policies should be able to allow or disallow certain programs, network packets, user activity, or other fields. However in order to make decisions on what to allow or disallow, the network administrator needs to know what is at the root of the problem they are trying to solve. This means that an optimal tool will allow a network administrator to track all network activity of individual endpoints or applications to enable examination and evaluation. This network tool should theoretically display to the network administrator what activity is normal, and what activity is unusual so they can most effectively create policies and rules for the network.

When reacting to an identified threat on a network, an administrator must have a tool that allows the quick localization of the threat, ability to isolate it, and eradicate it from the network. Reactive network administration relies heavily on a good alert management system that not only alerts the network admin of the presence of a threat, but does so before it has caused unwanted damage. Comodo Client uses its auto-containment tool to sandbox any program that has not been verified before in order to keep the program isolated until proven safe. An administrator could then be alerted to remove the program from a device. Stateful inspection can be used on larger enterprise networks to watch the activity of a specific endpoint on the network. When an endpoint's activity results in too many red flags, the endpoint itself can be blocked from connecting to the network.

Whitelists, blacklists, and behavioural inspection of traffic are all different methods that should be used by network administrators to ensure the safety of their network. Whitelisting is the act of allowing a behaviour that is known to be safe. In a purely whitelist approach, it is nearly impossible for malware or unintended applications to run on a network because they would need to be approved by the administrator first. Because of this process, this approach will successfully block zero day attacks and new exploits that have not already identified as malicious. The security of a purely whitelist approach is unmatched, but its practical implementation is where it has flaws. When all unidentified behaviours are disallowed by the network, there is a very large amount of activity blocked by the firewall that may not be malicious. This limits the users of the network to a finite list of actions possible to conduct on the network without seeking further approval by the network admin to add something new to the list. The job of the network administrator in this system is to manage the overhead of new requests and to make sure that the whitelist of acceptable network activity is as thorough as possible rather than reactively responding to threats that have already infiltrated the network.

Blacklists help a network administrator manage a network by blocking activity that is known to be bad. Unlike the whitelist approach, a purely blacklist approach will never block an activity that is not malicious, but in doing so, will allow activity that should not be present on the network. Although the purely

blacklist approach is much more friendly to the users on the network than a purely whitelist approach, it does not have the ability to block new threats that are not yet on the list.

The strategy that a network administrator chooses to use does not need to be used throughout the entire network. For example, a more whitelist approach may be more appropriate for managing programs and applications and a more blacklist approach may be more appropriate for managing HTTP traffic. In areas of the network with consistently known activity, a network admin can reap the security benefits from using a whitelist approach without the overhead of continuously updating a whitelist. In areas of the network that the origin or type of activity is always changing, a blacklist approach may be more suitable than a whitelist approach because a much smaller amount of the acceptable activity on the network will have been proactively identified as acceptable and therefore permitted in a whitelist system.

Because neither the purely blacklist approach or purely whitelist approach is optimal, a network administrator should use a combination of the two approaches, and develop a method to filter the activity that is not yet classified by either list. Heuristic algorithms and data analytics can help a network administrator attempt to profile activities that fall in between the scopes defined by blacklists and whitelists. These algorithms need to be able to identify when an endpoint or application is operating outside of what is perceived to be normal on the network, and they need to be able to either act on that abnormality or alert the network administrator to take action.

With the ubiquity of the modern internet, it is critical that data be protected from security threats. There exist countless forms of network security threats, both in the form of intentional and unintentional threats. Threats are ultimately broken down into one of four categories, defined by William Stallings in his book *Network and Internetwork Security: Principles and Practice*:

- Interruption - An asset of the system is destroyed or becomes unavailable or unusable
- Interception - An unauthorized party gains access to an asset
- Modification - An unauthorized party not only gains access to, but tampers with an asset
- Fabrication - An unauthorized party inserts counterfeit objects into the system

The field of network security can ultimately be summed as the field of preventing any of these attacks from happening as best as possible, and mitigating the damage an entity (we will specifically be focused on the field as it pertains to companies and other large organizations) will suffer in the event that one occurs. Network Security Analysts have a variety of tools that may be utilized in order to protect the data of their organization; for this study, we will be focusing on the use of Firewall technology.

In their study on the characteristics of Internet Background Radiation, Pang et al. noted that “perhaps the most striking result . . . is the extreme dynamism in many aspects of background radiation . . . the mix of background radiation sometimes changes on a nearly-daily basis. This dynamism results in a pot-pourri of connection-level behavior, packet payloads, and activity sessions seen in different regions of address space” [7]. These characteristics help bring to light the reason why firewalls can be very difficult to configure to properly allow all legitimate traffic and block all malicious traffic; network security is an arms race between hackers and analysts, with hackers attempting to mimic legitimate traffic in order to gain access to sensitive data. Take, for instance, the common Microsoft Office macro attack vector: a common means of attacking institutions that regularly accept large quantities of text documents, such as universities during application season, is through malicious code inserted into seemingly-benign Microsoft Word documents. For example word processors can open a network connection so hyperlinks like `www.google.com` will launch a connection from whatever application is used to view the document. Thus, a policy that simply blocks all network traffic initiated by Microsoft Word in an attempt to block macro attacks will also block legitimate traffic as a byproduct, potentially interrupting the workflow of members of an analyst’s organization. Correctly blocking all (or at least a significant majority) of malicious traffic requires a more nuanced approach.

2.4 Filtering by intention

As PEACE presents a network analyst with more insight into the intentions of network users, it makes sense to create network policy that filters based on intention. Filtering by intention refers to making decisions about the validity of traffic based on the reason why it was sent to the network rather than its content. This is advantageous because it helps a network analyst classify the grey traffic that does not appear in their black and whitelists.

In their study on detecting malware by tracking user intentions, Jeffrey Shirley and David Evans found that, compared to traditional access control, intention-focused access control offers easier policy development due to referencing “higher-level abstractions of intent rather than lower-level application behavior”, resulting in policies that are “more readily comprehensible to humans and amenable to automated policy development” [10]. They also found that intent-based policies “promote greater reusability of policies, since they do not depend on the specific details of how individual applications carry out that intent”. In addition, they found that intent-based policies are more usable as they require less extensive user configuration. This research suggests that not only are intent-based policies more readily understandable, but that they are also quicker and easier to develop by network analysts, potentially streamlining their workflow.

3 Methodology

In order to understand the value of the additional information provided by the PEACE system, we built a web application to present the data to participants. The creation of a web application was the best solution because of its ability to efficiently collect data, and its ability to fully guide the participants through the information we were presenting them. By creating a web application, we were able to standardize a user interface in order to fully isolate the value of the presentation of different types of network traffic data. The web application was responsible for giving the participant general information about the PEACE system and their role as a network analyst, walking the participant through a tutorial to teach them the layout of the web application and flow display, and finally walking them through each of the three phases of the study.

3.1 Designing the study

When designing the user study, we decided to conduct our study in three main phases. In order to isolate the impact of the PEACE data, it made sense to have separate segments for both the data sets that included PEACE data and those without. The three segments were derived from two different data sets of seven network flows, and were designed to best show how the user’s decision to block and allow network traffic changed as a result of the PEACE data. The first segment of the study displayed its dataset of network flows to the user with only the network action time, source IP address, destination IP address, destination port, source port, destination host, protocol, flags, and application path. The second segment presented a different set of network flows with the same surface level information as the first segment, but also incorporated PEACE-exclusive insights like keystrokes, mouse clicks, and GUI text. The third segment revisited the same data set as the first segment, but this time it included the PEACE-exclusive information described in the second segment.

In order to minimize systematic bias across trials, we randomized the order of segment presentation for the first two segments, and displayed the third segment last. This randomization made sure that some participants started phase one of the study with PEACE flows, and others did not. The ordering of the segments was important to the integrity of the results because of the possible influence of one segment on another. For example, it did not make sense to display a dataset with PEACE insight before the same dataset without the PEACE insight because of the possibility of the participant recalling information that should not be available to them to make their decision. This meant that we needed to make sure that within a specific dataset, the non-PEACE segment would be presented to the participant before the PEACE dataset. The order of the first two segments shown to the participant could be randomized because they were of different datasets. However we still chose to randomize those segments because we did not want the type of data first presented to the participant to affect their process of determining the validity of a network flow. For example if the first segment presented to every participant was a segment

with PEACE data, they may develop a pattern of study different than if the first segment presented was one without PEACE data. In order to minimize this influence, we randomized the presentation of segment one and segment two as described above. We labeled the order of the three segments as seen by the participant as phase one, phase 2, and phase 3. Phase one therefore consisted of seven flows with or without PEACE insight, phase two consisted of seven flows from a different dataset than phase one and had the opposite data type (PEACE vs no PEACE), and phase three consisted of the same seven flows previously presented without PEACE insight however this time presenting it with the PEACE insight.

3.2 Designing the flow page display

In order to get the best results from the user study, it was critical to optimize and standardize the display of the network flows. We first needed to decide how we wanted to display the 21 network flows (seven for each of the three phases) to the user. We initially decided to follow the standard and display all of the flows in a single table, but ultimately decided to present one flow at a time to the user.

The one flow at a time approach was more effective than one large table of information because of its ability to present data to our participants. Because of the independence of each network flow in our study, there was no need to display all of the flows on the screen at once. This approach makes more sense for a firewall application because there is too much network activity on a live network to look at each one individually. However, because this study focused more on the impact of the information provided by the PEACE system, it made sense to optimize the users ability to look through all of the information presented to them. By only displaying one flow at a time, the participant may have been able to make a more educated decision to block or allow that specific network flow with this design.

When designing the flow page, we needed to consider the best way to present the PEACE elements, the non-PEACE elements, and GUI elements. We chose to display the PEACE data in a table for clicks and keystrokes, hierarchical text for the GUI Text, and displayed the non-PEACE data in a grid at the top of the page. We wanted to make sure that the location of the flow elements would not change from flow to flow in order to allow the participant to build a consistent search pattern as they progressed throughout the study. By placing the non-PEACE data at the top of the screen, the beginning of the process for analyzing a network flow remained the same for both phases with and without peace data. This design can be seen in Figure 1 above.

When designing the flow page, the first element displayed to the user was the non PEACE related flow information. This segment was presented first to the participants because it is the only segment that is included in every flow regardless of phase. By placing this component first, users could begin to develop a pattern for analyzing the flow information because of the consistent layout of the display. This component was also laid out in a three by three grid

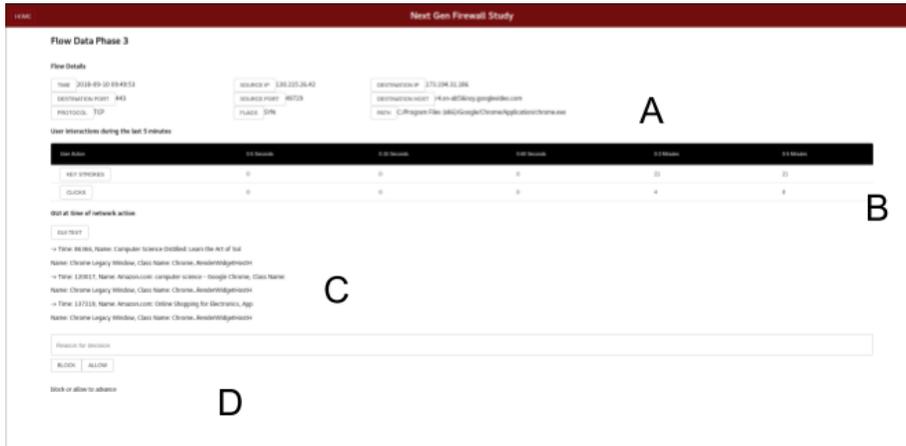


Figure 1: Flow Page (With PEACE) - This is the web page application's display of a single network flow. A: Non-PEACE Flow Elements. B: Click and Keyboard Elements. C: GUI Text Element. D: Participant Decision Element



Figure 2: Non-PEACE Flow Elements (General Flow Information) - This element shows the general information a network analyst would have using the current industry standard firewalls

User interactions during the last 5 minutes

User Action	0-5 Seconds	0-15 Seconds	0-60 Seconds	0-3 Minutes	0-5 Minutes
KEY STROKES	0	0	0	21	21
CLICKS	0	0	0	4	8

Figure 3: Click and Keyboard Elements - This element shows participant the actions that occurred on the user's machine the moments before the packet was sent to the network.

in order to optimize the participants ability to efficiently analyze the data. By spreading out the metrics as much as possible, the participant would get used to looking to different places of the screen to obtain different information. This is much easier for a user than scanning through a list or block of text to find a piece of information that they are looking for. For example, if a user wants to look for the PATH element of the data flow, they will become used to looking to the bottom right corner of this segment on the screen.

Each metric label of the general information segment was implemented as an indicator button. When a participant found a particular metric critical in their decision to block or allow a network flow, they were instructed to indicate that metric by clicking on its label. By making each metric label a button, we were able to better gather information about which metrics participants were using to make their decision. Compared to having the participant verbally articulate or write down their thoughts as they progressed through the study, having them select indicators allowed us to easily classify the influences of their decisions without needing to individually parse all text or audio recordings to understand the users thoughts. These buttons were also placed directly next to each piece of information so that the participants could click on them as they analyze each metric.

During the phases of the study that include PEACE elements, a mouse click and keyboard table is presented to the user. This information describes the number of keystrokes and mouse clicks as a function of time that occurred on the machine that the network flow originated on. The five different time windows that were presented to the participants were zero to five seconds, zero to fifteen seconds, zero to one minute, zero to three minutes, and zero to five minutes. It made sense to present all of these intervals to the user because it is how the PEACE system recorded both the mouse clicks and keystrokes. It is also important to note that these intervals are cumulative because it is more intuitive to think of the number of user interactions in between the time that the action occurred and the flow being received on the network rather than seemingly arbitrary windows of classification. For example in the Figure 3 above, one can see that there were 0 keystrokes or mouse clicks within a minute of the network action, 21 keystrokes and 4 mouse clicks within three minutes of

the network action, and 21 keystrokes and 8 mouse clicks within five minutes of the network action. It is also clear from this data that there were zero keystrokes between minutes three and five, but there were an additional 4 mouse clicks.

In addition to the user action table, during the PEACE phases of the study, the GUI Text hierarchy is presented to the participant. The GUI Text is the element of the PEACE data that describes in text what is shown on the user's screen at the time of the network action. The text is broken down into a hierarchy of different components such as Time, Name, Class Name, and Class Text. Each arrow character '->' followed by a time value indicates a new snapshot of the screen before the network action. For each snapshot the screen is described using classes. For example Figure 4 above shows that at time 6445 PowerPoint Presentation 1 is open and within that a workspace, a pane, a slide, a text box, and a hyperlink. This snapshot of the screen was displayed as such so that the participant could distinguish between the different levels of depth of the UI at the time of the user's interaction with the network. Notice this description of the user screen is different than at time 579 when the user is interacting with the NetUIToolMenu of PowerPoint because the user was interacting with a different component of the application at that time. Again like other metrics, we have provided a button for the participant to indicate if the GUI Text influenced their decision to block or allow the network flow.

The last segment of the display page is the panel for the participant to make their final decision to block or allow the network flow. We decided to place this segment at the bottom of every flow page because we wanted the participant to look at all of the data in front of them before making a decision. In this segment, the participant is prompted to type in why they made the decision that they did, and to decide whether or not to block the flow. We decided to use block and allow buttons for the decision in order to force the participant into making a decision meanwhile minimizing any potential bias of a default decision. Upon selecting to block or allow the flow, a "continue" button is presented to the user instead of the "block or allow to advance" message.

3.3 Building the web application

We decided to present our user study with a Nodejs-React web application because of its ability to let us fully customize the user experience during our trials and accurately collect data. Nodejs is a backend javascript framework that easily integrates with the third party Herokuapp[6], which we used to serve our web app [3]. React is a new client side framework created by Facebook designed for single page web applications [8]. We chose to design the web application using the React framework because of its ability to contain the entire functionality of the user study on the client side which improved usability and even accuracy of data collection.

The single page web application works by dynamically rendering individual components of the page when the state of that component is updated. We used this property of the React framework to make our web app render immediately upon the user's actions and record them rather than contacting an external

GUI at time of network action

GUI TEXT

-> Time: 579, Name: Context Menu, Class Name: NetUIToolWindow, Class Text: Menu

Name: Context Menu, Class Name: NetUIToolWindow, Class Text: Menu

Class Name: NetUIPanViewer, Class Text: Pane

Class Name: NetUITWMenuItemGroup, Class Text: Group

Name: Open Link, Class Name: NetUITWBtnMenuItem, Class Text: Menu Item

-> Time: 6445, Name: Presentation1 - PowerPoint, Class Name: PPTFrameClass, Class Text: w

Name: Workspace, Class Name: MDIClient, Class Text: pane

Class Text: pane

Name: Slide, Class Text: pane

Name: Slide 1, Class Text: Slide

Name: Subtitle TextBox, Class Text: textbox

Name: <http://myshreddies.com>, Class Text: hyperlink

-> Time: 9293, Name: Presentation1 - PowerPoint, Class Name: PPTFrameClass, Class Text: w

Name: MsoDockTop, Class Name: MsoCommandBarDock, Class Text: pane

Class Name: MsoCommandBar, Class Text: tool bar

Name: Ribbon, Class Name: MsoWorkPane, Class Text: pane

Figure 4: GUI(Graphic User Interface) Text Element - This element describes the user interface of the origin application at the time the packet was sent to the network

The image shows a user interface element for a decision. It consists of a large rectangular text input field with the placeholder text "Reason for decision". Below this field are two smaller, rounded rectangular buttons. The left button is labeled "BLOCK" and the right button is labeled "ALLOW". Both buttons have a thin border and a light gray background.

block or allow to advance

Figure 5: Participant Decision Element - This element shows how the participant would indicate their decision to block or allow a particular data flow

server during intermediate steps. When the state of the web page was updated.

The web application was broken down into a couple main React components, the `AppComponent`, `SingleFlowPageComponent`, `StartPage`, and the `PhaseBreakPage`. Upon serving the web app, Heroku served the proper components depending on the route specified in the URL. For example all calls to the `"/` directory and any route that included the `"/` (which is included in every route) would serve the `AppComponent`. This makes sense because the `AppComponent` was in charge of the the structure, navigation, and style of the webpage, which needed to be consistent across every page of the web app. If the route exactly matched `"/` then the participant would also be shown the `StartPage` component within the `AppComponent`.

Upon reading the background information of the study presented on the `StartPage`, and clicking on "Start Tutorial", the participant was finally brought to the `SingleFlowPageComponent`, where the bulk of the functional web app is handled. The participant's click on the "Start Tutorial" button calls for the `"/tutorial` endpoint of the application which returns the first flow of the study to the participant with a tutorial wrapper. In this situation, the `SingleFlowPageComponent` is passed the tutorial indicator upon construction in order to properly guide the user through the example PEACE flow and make sure that user actions are not yet stored as results. The tutorial indicator tells the `SingleFlowPageComponent` to render each element of the PEACE flow one at time, blurring out the others, and to wait for the participants indication that they are ready for the next step by clicking the "Next Step" button. The "Next Step" button merely changed the style of the HTML elements on the tutorial page in order to walk the participant through the tutorial until they were ready to start the study. This design choice for the tutorial was optimal because it allowed us to inject instructions and walk the participant through the `SingleFlowPageComponent` without modifying the functionality of the actual study and without rewriting much of the same code.

Upon completion of the tutorial and the click of the "Start Study" button at the end of the tutorial page, the participant was then sent to their first

flow page. At this point, the web app randomized which flow to display first and sent the participant to either “/flow0” or “/flow8”. When the “/flow:id” route was called, the `SingleFlowPageComponent` was constructed with the flow id as specified in the route, and access to the logs which are stored in the `AppComponent`. The `SingleFlowPageComponent` then retrieved the proper flow stored in a datafile and displayed it. The display type, PEACE versus non-PEACE, of the flow was dependent on the randomization and stored in the state of the `SingleFlowPageComponent`. As the participant navigated through the web app, interactions with the web app got pushed to the logs, and each time the participant classified a flow and was ready to move on the next one, the `flowid` parameter of the state was increased by one resulting in a dynamic re-rendering of the `SingleFlowPageComponent` with a new flow to display to the participant. We chose to utilize this type of flow navigation to take advantage of the strengths of the React Framework. By making changes to the state of the `SingleFlowPageComponent` and dynamically re-rendering the page, we were able to quickly show the participant new information without reloading the page and without contacting an external server.

At the end of each phase, the `SingleFlowPageComponent` would redirect the participant to the `PhaseBreakPage` before routing them to the first flow of the next phase when they were ready to proceed. After the participant had classified the last flow of the third and final phase, the `SingleFlowPageComponent` displayed a link to save and download all of the logs during the duration of the study.

3.4 Gathering network flows

In order to best identify the value of the additional information provided by the PEACE system, we chose network flows collected by the PEACE system to present to participants to analyze. A network flow is a single network interaction that has been received by the router and saved in the PEACE data system. At the lowest level, these network flows are what a network analyst or IT specialist would need to look over in order to fully secure the network. In the traditional firewall systems, only surface level information is available for the analyst to look through to make their decision to block, allow, or create policy on a specific flow, or type of data flow. Because the PEACE system also collects all of the surface level information presented to an analyst by a standard firewall, we were able to retrieve all of our data flows for the study from the PEACE database, and selectively display different depth of information to our participants.

When picking data flows from the PEACE database, it was important to chose dataflows that had the elements of PEACE data that we were interested in testing. Because not all network traffic originates within the network, not all of the flows in the PEACE database contain more than the surface level metrics for the analyst to make decisions from. This is because the PEACE software must be running on the machine in order to collect these additional metrics. When conducting the different phases of this user study, we wanted to see how the PEACE data influenced the users decision to block or allow the

flow, therefore it was necessary for most of our data to contain information on the users' graphic user interface at the time of the network action, and the count of the user's most recent clicks and keystrokes.

Additionally, we needed to make sure there was a wide variety in user actions that were being captured by the PEACE system. A wide variety of network flows is necessary because different types of network actions may appear differently on the PEACE system. PEACE is able to show how different applications interact with its users and with the network in different ways. For example a Windows defender automated process may not seem to have any GUI Text at all, but this is to be expected of a background process. Microsoft Excel may have a very extensive pane layout described by the GUI Text, and one might expect a mouse click to initiate any network requests from such an application. By looking through the PEACE database for different applications we were able to compile a wide range of network flows for the study.

Although we were able to compile a wide variety of network flows, we were not able to simulate an exhaustive list and therefore needed to manually create additional flows. In particular, we needed to create seemingly malicious data flows because they even if they were in the PEACE database, they were not readily identified as such. We needed to have malicious data flows as part of the user study in order to give the participants reason to block or allow network flow. When creating dataflows for the study, we were able to take components of flows from the original PEACE database and tweak them for a different purpose. For example, eliminating mouse clicks from a network flow that came from Microsoft Word could now indicate a malicious script running in the background of the application instead of a user clicking on a hyperlink. The majority of the other metrics in this situation could be left exactly the same.

3.5 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 user, 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. The purpose of conducting this study was to determine quantifiably how users interacted with the PEACE specific insight and whether or not it helped them make decisions.

At the beginning of each trial, the participant would be 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, the identity of Akamai as a known content delivery company was provided after feedback expressing suspicion of the name; care was taken to clarify to participants that the name "Akamai" being present within a network flow did not inherently make it either malicious or trustworthy, but rather that the name itself was not suspicious.

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.

Indicators Used on Non-PEACE Flows

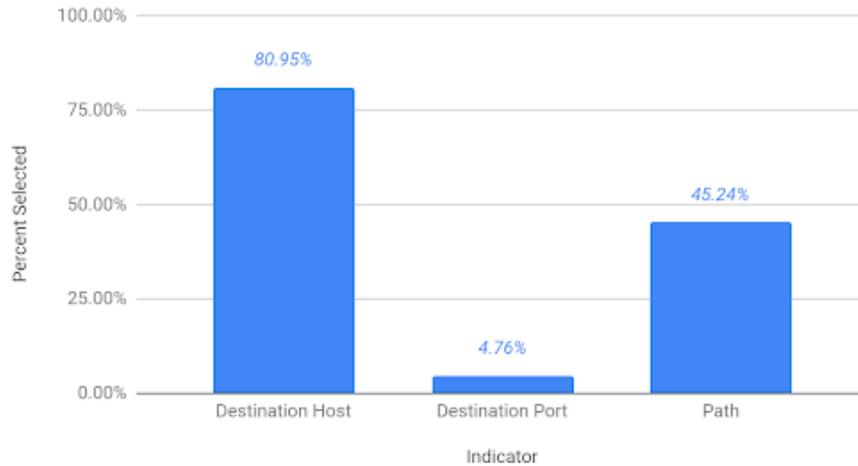


Figure 6: This is a graph displaying the likelihood of an indicator being selected by a participant for any non-PEACE flow

4 Results

Once the trials with each of the six participants were completed, the decisions and actions of the participants during the study collected by the web application were compiled and analyzed. This result section will identify how the participants of the study interacted with the different network flows, and how begin to illustrate how the PEACE exclusive information changed the decision making process of all participants in the study.

4.1 Indicator utilization

In Figure 6, it is clear that the destination host was the most widely used indicator for any non-PEACE data flow having been utilized on 80.95% of decisions to block or allow. The application path was used 45.24% of the time, and the destination port was used to make 4.76% of decisions. All other basic firewalling metrics displayed in the non-PEACE segment of this study such as network action time, source IP address, destination IP address, source port, protocol, and flags were never used by any participant during the non-PEACE segment of the study to make a decision. These values were obtained by counting each time an indicator was selected by each participant for each of the seven presented non-PEACE network flows, and dividing that total accumulation of each indicator by the number of flows and the number of participants. These values are useful because they show the aggregate decision making tendencies of the participants

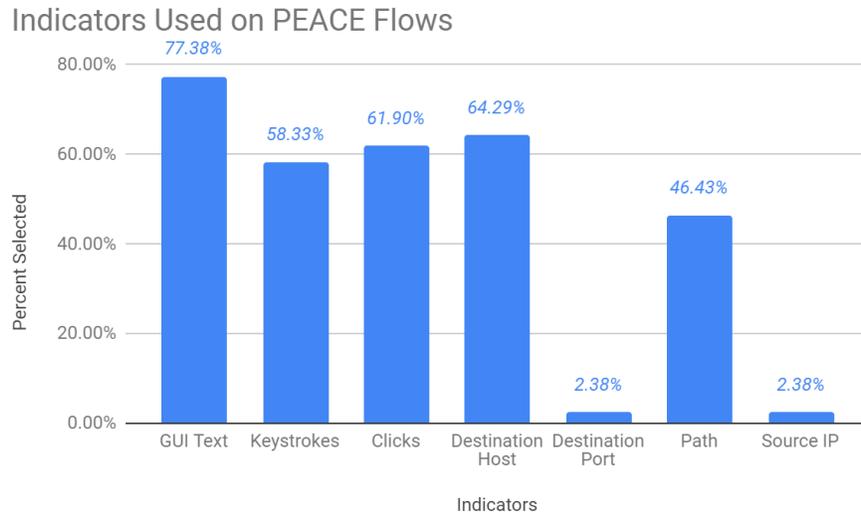


Figure 7: This is a graph displaying the likelihood of an indicator being selected by a participant for any flow with additional PEACE insight

when PEACE insight is not available to them.

In Figure 7, it is clear that the PEACE-exclusive insight was frequently used by the participants to make their decision to block or allow a network flow. Compared to Figure 6 above, this graph introduces the three new indicators of GUI text (77.38% utilization), keystrokes (58.33% utilization), and mouse clicks (61.90% utilization). The destination host was utilized 64.29% of the time with PEACE insight compared to the 80.95% without. All other non-PEACE metrics were utilized within a 3% rate during the PEACE included segments of the study compared to the those without PEACE insight.

4.2 Blocked flows for the non-PEACE phase by participant

Figure 8 helps visualize the decisions each of the participants made to block or allow each of the seven flows presented without PEACE insight. Each of the participants chose to block between two and four of the seven flows they were presented, and each flow was blocked between two and five times throughout the entirety of the trials. Of the 42 decisions to be made during the study, the participants cumulatively chose block 20 twenty times, and allow 22 times.

Flow	Participant						Sum
	1	2	3	4	5	6	
1	0	1	0	0	1	0	2
2	1	0	1	1	1	1	5
3	1	1	1	0	0	0	3
4	0	1	0	1	1	1	4
5	0	0	1	0	1	0	2
6	1	0	0	0	0	1	2
7	1	0	0	0	0	1	2
Sum	4	3	3	2	4	4	20

Figure 8: This table shows the decisions made by the participants for each flow in the Non-PEACE Phase of the study. A red one indicates the flow was blocked by the participant, and a green zero indicates the flow was allowed by the participant

Flow	Participant						Sum
	1	2	3	4	5	6	
1	1	0	1	0	0	0	2
2	1	0	0	0	0	1	2
3	0	0	0	0	0	1	1
4	0	1	0	1	1	1	4
5	1	0	0	0	0	0	1
6	1	0	0	0	0	0	1
7	0	0	0	0	0	0	0
Sum	4	1	1	1	1	3	11

Figure 9: This table shows the decisions made by the participants for each flow in the PEACE Phase of the study. A red one indicates the flow was blocked by the participant, and a green zero indicates the flow was allowed by the participant.

Flow	Participant						Sum
	1	2	3	4	5	6	
1	1	1	1	0	1	0	4
2	0	0	1	1	1	0	3
3	1	1	1	0	0	1	4
4	0	0	0	0	0	0	0
5	0	0	1	0	1	0	2
6	0	0	0	0	0	0	0
7	0	0	0	0	0	1	1
Sum	2	2	4	1	3	2	14

Figure 10: This chart describes the number of switches, or decision changes, each participant encountered when presented with flows from the same data set but different levels of insight

4.3 Blocked flows after introduction of PEACE insight by participant

Figure 9 helps visualize the decisions each of the participants made to block or allow each of the seven flows presented once PEACE insight was introduced. Four of the participants only blocked one of the seven flows in this phase of the trial, and the other two blocked three and four respectively. Flow seven was the only flow to be allowed by all six participants. As a group, there were fewer decisions to block a flow once PEACE was introduced compared to the initial phase resulting in 11 decisions to blocks and 31 decisions to allow.

4.4 Decision changes for each flow by participant

Figure 10 quantifies how the PEACE insight influenced the verdict reached by the participants during the study. Of the total 42 decisions made by the participants during the non-PEACE phase of this study (six participants were presented with seven flows each for a combined 42 decisions), 14 of them were made differently after the participant was presented with the PEACE insight. Every participant changed their mind at least once, and participant three changed their mind on four of the seven flows. Flow numbers four and six were the only flows that did not experience a change in verdict by any of the participants throughout the entirety of the study.

4.5 Direction of decision change for each switch

In the Figure 11, it is clear that the majority of the changed decisions were in the direction of block to allow. Throughout the entirety of the study, participants decided to allow 11 flows that they previously blocked, and only chose to block three flows that they originally chose to allow. Both flows four and six

Flow	Allow to Block	Block to Allow
1	2	2
2	0	3
3	1	3
4	0	0
5	0	2
6	0	0
7	0	1
Grand Total	3	11

Figure 11: This chart shows the direction of decision change for each of the 14 switches

never experienced a switch, whereas flows one and three experienced the most volatility having been changed four times each by the six participants.

5 Discussion

This section will describe the impact of the results explained in section 4 of this paper, and what is to be learned from this study. This study ultimately suggests that PEACE specific metrics had a large impact on the decision making process of the six participants because of the rate at which the metrics were utilized and because of the amount of decisions they changed once presented with additional insight. Although the results of this study very strongly suggest this, further studies still need to be conducted in order to determine the significance and impact of the depth of information presented by the PEACE system.

5.1 PEACE indicator utilization

During this study, all six participants incorporated the PEACE metrics in their decisions when available. When presented with all metrics, the GUI text was the most utilized indicator followed by destination host, mouse clicks, keystrokes, application path, destination port, and source IP address. The three PEACE-exclusive metrics examined in this study were among the four most utilized indicators across the board, all of which surpassing eight of the nine metrics used by the current industry standard. The destination host, which was the most utilized indicator during the phase without PEACE insight at 80.95%, was used 15% less frequently in phases when participants were able to rely on the GUI text, mouse clicks, and keystrokes to make their decisions.

The PEACE insight clearly had an impact on the participants of this trial, but this study cannot conclusively project an impact on the network industry because of the limited size of the study and the level of expertise of the participants. Because none of the participants were formally trained in network security, it is possible that they relied on the PEACE insight to make their decisions because of its relative ease of use. Metrics like keystrokes and mouse clicks are potentially easier for one with a non technical background to make conclusions from than a port number or protocol.

Although these metrics may be utilized differently by those with more experience, it is still evident from this study that the PEACE insight may make network security more possible for those with less experience. Even if these PEACE-exclusive metrics do not help the advanced network security analyst, this study suggests that for those of less expertise it does make a difference. The introduction of PEACE to the security industry could lower the barrier of entry if network traffic were more intuitive to categorize with the inclusion of GUI text, mouse clicks, and keystrokes.

5.2 Switched verdicts as a result of PEACE

The insight provided by PEACE had a large impact on the decision making process of the six participants because of the number of times they changed their original verdict once presented with the addition information. Every participant changed their decision to block or allow a flow at least once if not multiple

times during study. Not only were the GUI text, mouse clicks, and keystrokes incorporated in the decision making process of the participants, they influenced them to change their verdict. Participants collectively changed 33% of their decisions when presented with PEACE insight. This number could largely be skewed by the inexperience or lack of confidence of each of the participants original decisions, but it unquestionably had an impact.

Not only were the verdicts influenced by the PEACE insight, they were more often than not in the direction of block to allow. Of the 14 total switches, 11 of them were allowed to pass after initially being blocked. It seems as though the metrics provided by the PEACE system allow a network analyst to be more confident in the intent of the flow, and therefore are more willing to permit something that originally looked suspicious.

5.3 Considerations for future studies

Although this study gives valuable insight into the usefulness of information provided by the PEACE system, there is still room for research on this topic. The results of the study were particularly limited by the background of the participants, and the flows they were presented during the trial. Future studies should evaluate the effectiveness of the PEACE-exclusive data among trained network professionals. Due to the scope of this study, none of the participants were formally trained in network security. A trained network professional therefore may use the PEACE information in a different way than the participants of this study did. For example, a trained network professional may be more likely to correctly interpret the intent of a network flow based entirely on its general characteristics. Therefore the effectiveness of the PEACE information could be less pronounced in a more experienced pool of participants.

Future studies should also be conducted with more concrete network flows. When choosing flows from the PEACE database, the exact origin was unknown. This made it inherently difficult to know if a network flow from the database was malicious or not. Because of this uncertainty, this study was unable to measure the participant's accuracy when classifying network flows. By performing a variety of network actions on a PEACE monitored system and collecting data, a future study could have a full understanding of the origin of each network flows, and therefore could make more definitive claims about the success of the participant's flow categorizations and ultimately the effectiveness of the PEACE-exclusive information.

5.4 Conclusions

The study is ultimately suggestive, rather than conclusive, due to the small sample size of six participants and their technological background. Despite this, the preliminary results strongly suggest that the metrics provided by PEACE did have a large influence on the participants decision making process; participants 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.

References

- [1] Alam, M. Afshar., et al. Proceedings of National Conference on Recent Developments in Computing and Its Applications, August 12-13, 2009 . I.K. International Pub. House, 2009.
- [2] Cheswick, William R., Steven M. Bellovin, and Aviel D. Rubin. Internet Security . Addison Wesley, 1996. Manum, Marcus. “Firewall Toolkit V1.0 Release.” Firewall Toolkit V1.0 Release , 21 Oct. 1993.
- [3] Cloud Application Platform. Heroku. www.herokuapp.com Accessed 3 November 2018.
- [4] Key PEACE Design Concepts. ContexSure Networks, 2016, www.contexsure.com, Accessed 10 October 2018.
- [5] Next Generation Firewalls. Palo Alto, 2019, www.paloaltonetworks.com/products/secure-the-network/next-generation-firewall, Accessed 8 October 2019.
- [6] Nodejs. Joyent Inc., www.nodejs.org, Accessed 3 November 2018.
- [7] Pang, Ruoming, et al. “Characteristics of Internet Background Radiation.” Proceedings of the 4th ACM SIGCOMM Conference on Internet Measurement - IMC 04 , 2004, doi:10.1145/1028788.1028794.
- [8] React - A Javascript library for building user interfaces, Facebook, www.reactjs.org, Accessed 3 November 2018.
- [9] Roesner, Franziska, et al. “User-Driven Access Control: Rethinking Permission Granting in Modern Operating Systems.” 2012 IEEE Symposium on Security and Privacy , 2012, doi:10.1109/sp.2012.24.
- [10] Shirley, Jeffrey, and David Evans. “The User Is Not the Enemy.” Proceedings of the 2008 Workshop on New Security Paradigms - NSPW 08 , 2008, doi:10.1145/1595676.1595683.
- [11] Stallings, William. Network and Internetwork Security: Principles and Practice . Englewood Cliffs, N.J., 1995. ACM Digital Library .
- [12] Timberg, Craig. “A History of Internet Security.” The Washington Post , WP Company, 30 May 2015, www.washingtonpost.com/graphics/national/security-of-the-internet/history/.