ABSTRACT

Title of Thesis: REAL-TIME CYBERSECURITY SITUATION AWARENESS THROUGH A USER-CENTERED NETWORK SECURITY VISUALIZATION

Kaitlyn DeValk
Master of Science, 2022

Thesis Directed by: Professor Niklas Elmqvist
College of Information Studies

One of the most common problems amongst cybersecurity defenders is lack of network visibility, leading to decreased situation awareness and overlooked indicators of compromise. This presents an opportunity for the use of information visualization in the field of cybersecurity. Prior research has looked at applying visual analytics to computer network defense, which has led to the development of visualizations for a variety of use cases in the security field. However, many of these visualizations do not consider user needs and requirements or require some predetermined user knowledge about the network to create the visuals, leading to low adoption in practice. With this in mind, I took a bottom-up, user-centered approach using interviews to gather user-desired components for the design, development, and evaluation of a network security visualization tool, called Riverside.

I designed a visualization that attempts to balance providing a comprehensive view of an environment while supplying details-on-demand. Riverside's key contribution is a data-driven,

dynamic view of a network's security state over time, meant to supplement an analyst's real-time situation awareness of their network. Riverside's system automatically partitions internal from external network components to visualize potential attack vectors across the entire environment. This research supports the need for further incorporation of users into the cybersecurity visualization development lifecycle. I call attention to key requirements for creating effective cybersecurity visualizations and specific use cases where visualizations can be leveraged to augment operational cybersecurity capabilities.

REAL-TIME CYBERSECURITY SITUATION AWARENESS THROUGH A USER-CENTERED NETWORK SECURITY VISUALIZATION

by

Kaitlyn St. Thomas DeValk

Thesis submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Master of Science
2022

Advisory Committee:
Professor Niklas Elmqvist, Chair/Advisor
Professor Michel Cukier
Assistant Professor Leo Zhicheng Liu

Acknowledgments

I would first like to thank my advisor, Professor Niklas Elmqvist, for his guidance not only on my thesis research but throughout my graduate school experience. You've been instrumental in my ability to successfully conduct research in the visualization field, and I want to thank you for guiding me throughout this process. Additionally, I'd to thank Professor Michelle Mazurek for her assistance in navigating the realm of usable security, as it was her class in my first semester of graduate school that ignited my passion for exploring this area of research. I'd also like to note my appreciation for all of the participants who made this research possible and my colleagues within the security field who provided invaluable support and feedback. I want to thank my partner, John, who has constantly supported my aspirations and was always there to encourage me through my struggles and setbacks. Last, I would like to thank my committee members, Professors Michel Cukier and Leo Zhicheng Liu, for their flexibility and willingness to support the culmination of my master's thesis.

Thank you to everyone at the University of Maryland whom I've worked with over the past two years! I wish you luck in your future endeavors, and thank you for any assistance you may have provided me along the way. This research was supported by the U.S. Coast Guard Advanced Education program.

Table of Contents

List of Tables

List of Figures

List of Abbreviations

CND	Computer Network Defense
COA	Course of Action
CTA	Cognitive Task Analysis
C2	Command and Control
EDR	Endpoint Detection and Response
MCA	Malicious cyber activity
NOC	Network Operations Center
SOC	Security Operations Center
IA	Information Assurance
IDS	Intrusion Detection System
IOC	Indicator of Compromise
IoT	Internet of Things
IPS	Intrusion Prevention System
IR	Incident Response
IT	Information Technology
SA	Situation Awareness
SIEM	Security Information and Event Management

SOTA	State of the Art
TTPs	Tactics, Techniques, and Procedures
VizSec	IEEE Symposium on Visualization for Cyber Security
WWHF	Wild West Hackin' Fest

Chapter 1: Introduction

Cybersecurity analysts are flooded with hundreds of security alerts on a frequent basis, which can lead to missed indicators of compromise (IOCs) or add to the difficulty of rebuilding timelines during incident response (IR). Situation awareness (SA) in cybersecurity has been highlighted as an area that requires improvement [1], and I take direct motivation from Kokulu et al. who interviewed Security Operations Center (SOC) employees to highlight the underlying issues of SOC operations. The most common issue that they noted from their participant responses was "low visibility into the network infrastructure and endpoints" [2]. A common sentiment shared by another interview study was that "[analysts] don't know enough about [their] internal environment to know what someone would be after" [3]. Additionally, some of the key capability gaps identified by these cyber defense analysts were a lack of effective tools for detecting "Privilege Escalation, Defense Evasion, and Lateral Movement" [3].

One of the original models for defining SA in dynamic environments is Endsley's 3-Level model consisting of perception, comprehension, and projection [4], but a more applicable term for the security field is "cybersecurity situation awareness," or cyber-SA. The term is used to describe situation awareness as it applies to cybersecurity [5] and has been defined as "awareness of any kind of suspicious [or] interesting activity taking place in cyberspace" [6] and "[the ability to] identify adversaries, estimate impact of attacks, evaluate risks, understand situations and make

sound decisions on how best to protect valued assets swiftly and accurately" [7]. In this thesis, I will use cyber-SA to refer to an analyst's ability to **identify** anomalous or suspicious activity within an environment, **analyze** the information present, and **recommend** appropriate actions, correlating to three levels of the traditional SA model.

Information visualization is a method for "generating graphical representations of information" to provide users meaningful insights from their data [8]. For this reason, visualization has been presented as a technique that can help enhance cyber-SA [9–12], leading to the development of a multitude of cybersecurity visualizations over the years [13]. However, user input is often ignored or never sought out during the design and evaluation phases [14]. Additionally, cyber defenders are inherently skeptical of "automated reasoning about their data in general," and thus tend to place mistrust in cybersecurity visualizations [15]. In order to cultivate widespread adoption of tools, user needs and requirements must be taken into account, and end users need to be involved in the development lifecycle.

However, human-centered approaches are not common in the field of cybersecurity visualization [1, 16]. Additionally, a survey from 2014 revealed that out of 130 papers from the past 10 years of the IEEE Symposium on Visualization for Cyber Security (VizSec), 46% of papers included no user evaluation component [17]. Furthermore, many of these security visualization tools try to take the decision-making out of the analyst's control or require manual intervention from the analyst to produce a visualization at all. Unfortunately, users oftentimes lack the expertise to create effective visualizations or find it a "time sink" [18], resulting in many of these visualization tools never being used in practice.

1.1 Proposed Solution

While attempts have been made to build all-encompassing network security analysis tools, the concept of defense-in-breadth [19] is lost. It's been shown that analysts need "a tool box of different visualizations tailored to the specific requirements of each category of use of visualization" [9]. Furthermore, cybersecurity analysts are constantly presented with various dashboards and visuals, but few visualizations provide a concrete view that balances abstraction with detailed insights. In this thesis, I attempt to see how visualizations can be designed to aid the security workforce in maintaining cyber-SA of their environments. I first analyzed past and current capabilities that security personnel use when performing their roles. I then carried out a user-centered development lifecycle for a network security visualization tool. Riverside is a data-driven, dynamic network security visualization tool that displays changes in network topology over time. Last, I conducted a formative usability study to benchmark my visualization system's ability for assisting security personnel in maintaining cyber-SA of their environment and make recommendations for future work in this area.

I built off previous user-centered visualization design research [14, 20, 21] by conducting 24 semi-structured interviews with network and security professionals to discuss their experiences. Interviews were used to garner information about what mechanisms participants use for maintaining SA on their networks and what, if any, mechanisms they thought would be useful in a network visualization tool. I performed qualitative coding and thematic analysis to determine several themes surrounding participant's capability preferences, industry truths, challenges with current capabilities, and the applications of information visualization to participants professions. Additionally, I coded technical visualization features and mapped them to low-level

actions that could be used to infer higher level analysis tasks [22]. Using the interview data, I developed a data-driven network security visualization that allows analysts to make informed decisions about the state of their network. I lean on state of the art (SOTA) work for visualizing dynamic graphs [23] and apply it to the field of cybersecurity.

Unlike other network security visualizations, Riverside does not require any manual tuning or user-provided information and immediately yields valuable network insights upon deployment. Riverside displays a recognizable and animated network view connected to a timeline for temporal navigation, providing dynamic situation awareness over time. To evaluate Riverside, I conducted a usability study [24] with 10 security professionals to analyze participant's achieved cyber-SA using Riverside. I followed the guidelines presented by Freitas et al. [25] to directly incorporate real users into my evaluation process and identify Riverside's primary use cases for providing general cyber-SA. I developed a network attack scenario where participants were tasked with identifying individual components of the attack chain, and finally, explain their holistic analysis of what occurred on the network to produce suitable response measures. I completed quantitative and qualitative analysis to measure participant trial times, completion rates, usability issues, and participant's perceived levels of cyber-SA. Additionally, I collected feedback using three sets of 5-point Likert scale questions surrounding participant's ease of use, experience, and efficiency with Riverside. I also collected verbal and written feedback from participants at the conclusion of their session.

1.2 Contributions

Thesis Statement. *I contend that for visualizations to effectively provide cybersecurity situation awareness, a two-pronged "bottom-up" development approach should be used: 1) the technical design components should be sourced from end users, and 2) the visualization needs to inherently display all facets of an environment without user intervention.*

I use the word "environment" here to refer to any aspect that directly impacts the SA of a network, which includes internal-internal and external-internal analysis components. My direct contributions in this thesis are interviews with 24 network and security professionals to gain insights from their experiences and apply them to the development of effective cybersecurity visualizations. I provide 14 themes to help guide cybersecurity visualization development as well as 24 technical visualization features mapped to low-level actions to assist with security analysis. I present a bottom-up development approach to build a dynamic network security visualization tool, Riverside, and evaluate it with a scenario-based usability study. Finally, I make recommendations for future work in the cybersecurity information visualization field and show how effective cybersecurity visualizations should be developed based on end user trials and feedback.

This research was presented as a work in progress at the VizSec 2022 Poster Session [26] and Wild West Hackin' Fest (WWHF) 2022 Toolshed Presentations [27] to gather additional feedback from the community:

> Kaitlyn DeValk and Niklas Elmqvist. Riverside: Dynamic Visualization of Network Traffic for Situation Awareness in Computer Security. In *Proceedings of 2022 IEEE Symposium on Visualization for Cyber Security*, Oct 2022. Poster presentation at

2022 VizSec Conference.

Kaitlyn DeValk. Riverside: A Network Security Visualization Tool. Toolshed Presentation at 2022 Wild West Hackin' Fest Conference.

I contributed the initial research proposal, conducted all of the interviews and usability trials, and developed the bulk of Riverside. The research was conducted collaboratively with my advisor throughout my master's thesis who provided secondary analysis at various stages in the research, as well as direction for the development of Riverside. The poster was presented by Dr. Elmqvist at the IEEE VIS 2022 Poster Session as part of the VizSec Conference, and it showcased the interview portion of this research along with some initial designs for Riverside. I presented the Riverside prototype as a conference presentation at the Wild West Hackin' Fest, a security-oriented community conference.

1.3 Outline of Thesis

In Chapter 2, I walk through the related literature across four areas to include discussions with network and security professionals, visual analytics applied to cybersecurity defense, network security visualization tools, and user-centered cybersecurity visualization. I additionally provide a comparison taxonomy of similar cybersecurity visualizations to showcase where Riverside differs from other tools. In Chapter 3, I discuss the results of my semi-structured interviews with 24 network and security professionals. In Chapter 4, I walk through the design and development of Riverside. In Chapter 5, I review the results of the 10 usability trials for Riverside through a scenario-based evaluation. Chapter 6 provides the conclusion to my thesis and introduces future work.

Chapter 2: Literature Review

Prior work that is relevant to my thesis is spread across four areas: discussions with security professionals about their roles and experiences, applications of visual analytics in cybersecurity, cybersecurity visualization tools, and user-centered techniques for cybersecurity visualizations.

2.1 Discussions with IT and Security Professionals

One of the largest issues in the field of usable security is lack of feedback [28], but input from users is imperative to ensure that usable and effective security mechanisms are developed. To do this, researchers employ a variety of methods to understand the challenges that security professionals face to build tools that will truly help them. The most common are interviews which have been used to understand the challenges system administrators face in keeping systems or networks updated [29] or identify the challenges security professionals who assist in vulnerability discovery and the constraints that they feel prevents them from being successful in their roles [30]. While these interviews were not conducted with the same use case as mine, they were performed with the goal of understanding the obstacles that professionals face in their roles with the aim of developing effective solutions for them.

My interviews included professionals in the fields of network administration, SOCs, and Network Operations Centers (NOCs), as those are all professionals who could use network-based

visuals to assist in their job tasks. Similar studies have worked with SOC analysts or network and system administrators to understand the issues they face in their professional roles. Goodall et al. performed in-situ interviews with intrusion detection analysts and found that collaboration in organizations and the community was the primary issue for these analysts [31], while Kokulu et al. looked at the issues surrounding SOCs, both technically and culturally, to determine that SOC personnel's most common challenge was low visibility on their networks through interviews with analysts and managers [2]. This key finding of low situation awareness was echoed in NOC environments through a various in-situ user studies conducted by Paul [32].

Souza et al. conducted a survey with system administrators to understand their primary needs and found that not only did the administrators hold a variety of job responsibilities, but the largest problems they faced were "information quality and dynamic knowledge management" [33]. Another study used an ethnographic approach to understand the wide range of tools used by IT security professionals to develop better tooling and interfaces. They found that people prefer to use a handful of tools compared to just one or many, and the main issue with current tooling was "tailorability," or the ability to modify a tool based on an analyst's needs [34].

2.2 Visual Analytics Applications to Cybersecurity

Several papers have focused on understanding cybersecurity analyst's SA mental models using interviews [5], interactive tasks [35], and cognitive task analysis (CTA) [36,37], but all with the goal of determining how visual analytics can be used to developed effective cybersecurity visualizations by understanding the needs and workflows of analysts. Many of these studies focus on the use cases or applications of visualizations, such as D'Amico et al. who interviewed

cybersecurity professionals to confirm or deny assertions regarding the jobs of defensive cyber personnel to garner how they felt about security visualizations [18]. This built off previous work that focused on using CTAs to define cybersecurity defense roles, analysis, and workflows to produce recommendations for designing effective cybersecurity visualizations [38]. Many of study's findings were in line with what I uncovered during my interviews. However, where they showcase very broad results, my analysis was focused on gathering specific technical features that participants desired beyond the general use cases of cybersecurity visualizations.

Lavigne and Gouin explored how visual analytics can be applied to cybersecurity by categorizing different types of visualizations and their applicability in security [11]. Tyworth et al. looked at modifying Endsley's traditional 3-level SA model to better fit cyberseurity since "cyber-SA lacked well-defined boundaries but required collaboration across multiple entities" [12]. D'Amico & Kocka expanded Endsley's stages of SA and connected them to the stages of information assurance (IA) analysis based on possible use cases. They found that no one visualization can possibly cover all the stages of IA analysis or levels of SA and that there is no "silver bullet" for IA visualizations [9]. I build off this prior work by developing a network security visualization tool that is meant to provide high level insights to augment a security analyst's capabilities.

2.3 Cybersecurity Visualizations

The use of dynamic network graphs is not new [23] but applying them to cybersecurity can be a difficult task, as there is a lot of information available and a wide range of tasks that security analysts must perform. A few papers have completed surveys of pre-existing network security tools to showcase the wide range of cybersecurity visualization use cases. Shiravi et al. per-

formed a survey of network visualizations, categorizing them by use cases, such as port activity and internal-external monitoring, to emphasize the importance of real-time analysis, user experience, and the importance of usability evaluations [39]. Guimarães et al. built on previous work by conducting a thorough literature review and categorized 285 papers surrounding network and security visualizations by taxonomy in security and information visualization. They found that areas such as software-defined networking, Internet of Things (IoT), and human-centered evaluations are underexplored in the security information visualization field [16]. Other researchers have explored how general network visualization techniques can be applied to the security field to create more effective visualizations and aid analysis. They found that many of the current security visualizations lack "scalability of their visual metaphors, explicit representations of network topology, and heterogeneous network data" [40]. Du et al. proposed a framework for executing real-time, big data network security visualizations [41] aimed at combining various sources of network data and effectively visualizing the dynamic network traffic with a geographic visualization prototype. Riverside attempts to tackle some of these challenges while presenting an effective visual of a network's security state over time.

There are a multitude of cybersecurity visualizations out there with varying focuses, typically dependent on their the data source. These range from packet capture visualizations [42,43], firewall logs [44,45], IDS alerts [46,47], and IoT network traffic [48]. Other visualization tools were built with specific specific use cases in mind such as visualizing vulnerabilities for security assessments [49,50], impact to business and network operations [51,52], port security [53], or detecting anomalous traffic in large IP spaces [54–56]. NAVSEC [57], DAEDALUS-VIZ [58], and CyberVis [52] use 3D visualizations to display a variety of network security aspects with the goal of providing a more tangible way for analysts to visualize and "navigate" a network. Isis

used linked timeline and histogram event plots of network flow data to support analysis pivoting between various use cases [46].

Ocelot [21] uses mechanisms to Riverside, but it focuses primarily on grouping internal nodes by quarantine groups and uses sliding time windows whereas Riverside uses temporal navigation through a timeline to display past and current network insights. Additionally, Ocelot's layout uses circle packing combined with node-link diagrams where the remote, or external hosts, are placed in a circular layout, while Riverside exclusively uses node-link diagrams on an infinite canvas. NetCapVis [42] and VIAssist also use a client-server architecture with data batching for real-time analysis, but the former focuses on visualizing packet captures, similar to Wireshark [59], whereas the latter uses geographic chart visuals.

Situ also uses network flow data for a portion of its security dashboard visualization but provides a very rudimentary network view with little user customization and focuses more on anomaly detection using machine learning techniques [47]. Other visualization layouts that use timeline views similar to Riverside only allows users to choose various time spans [60]. The Time-Based Network Traffic Visualizer (TNV) focuses on visualizing network traffic over time with a timeline axis per host in a matrix plot [61]. TNV displays the links between hosts across timelines, but it requires the analyst to choose the data they're visualizing, whereas Riverside provides this feature automatically for network hosts.

2.3.1 Cybersecurity Situation Awareness Visualizations

There are several visualization tools that focus specifically on the use case of providing SA, akin to Riverside. Some of the first tools to tackle the problem of SA in security were NVi-

sionIP [62], VisFlowConnect [63], VISUAL [64], and FloVis [65], which set the stage for later visualization tools and all relied on network flow data. NVisionIP was one of the first dynamic network security tools and uses three views to show the data in a "galaxy view, small multiple view, and machine view" through scatter plots and bar graphs. VISUAL and VisFlowConnect use a grid layout with lines connecting the visualization entities, where FloVis uses a mix of radial edge bundling, 3D, and grid layouts to display the network data. Another tool called IP Matrix used matrices to visualize cyber attacks by using pixels to represent sites and colors to represent attacks [66]. OverFlow was a visualization tool that built off of the previous work of FloVis to create a central view that uses a concentric circle layout of node groupings to represent different segments of the network connected by links to represent communication [67]. Unfortunately, many of these tools were built off technology that has not withstood the test of time with the requirement of client-based software or the use of unmaintained frameworks and tool-kits.

VisAlert [68] was built for the use case of SA but exclusively visualized IDS alerts in a hierarchical circle layout with a node-link diagram of the network in the center. It focused on providing analysts insight in network intrusion detection but required the user to manually aggregate their data for input into the tool. Similar to VisAlert, Zhou et al. used a radial layout with edge bundling and node-links to create NetSecRadar but focused on tackling network data fusion while providing real-time event correlation to users [69]. Best et al. developed two tools to create a real-time, SA platform called MeDICi [70]. Their tools CLIQUE and TrafficCircle combined to provide network behavior graphs presented in a grid-row layout and a radial "time wheel" to convey network communication data, respectively. Erbacher developed a SA dashboard that used circles, or "gauges," to represent a system and was aimed at assessing mission impact for decision-makers, which involved risk and vulnerability scoring as well as evaluating network

security components [71]. Similarly, Dagger was proposed as a way to model and visualize mission through hierarchical layering techniques like the sunburst visualization [72].

BANKSAFE uses a combination of circles to represent 24-hour time-series data along with treemaps and matrix grids for activity data all by individual hosts through a hierarchical mapping of organizational levels and policies [73]. OCEANS [74] and CyberSAVI [75] were focused on creating collaborative security dashboards for overall SA, but the former used various graphical charts where the latter presented a high-level topological view of a network using node-link diagrams. Despite many of these tools providing great insights into a user's network, few of them discuss real-world deployment, or they contain complicated graphics and dashboards that make it hard for users to immediately comprehend.

Some more modern SA visualization tools include NStreamAware [76] and CyberPetri [77] that both provide real-time network security analysis. NStreamAware uses "time slices" and data streams to collect and display relevant network security information while CyberPetri is a later rendition of Ocelot [21] that was re-designed for monitoring a network security competition. CyberPetri used 15 minute data batching periods and colors to encode specific events, while Riverside uses colors to differentiate different components and batches data every two seconds allowing more accurate temporal analysis. NStreamAware uses data streaming and provides a multitude of dashboards to show slices of data over time, compared to Riverside which is focused on providing a concrete network environment view over time. Gray et al. built a network topology visualization tool, but it focused on visualizing external entities and providing SA to network administrators of potential external security threats [78]. Some of these tools and frameworks incorporated user-driven feedback and requirements through results of other studies or methods of their own, but they all use varying layouts and features to provide SA. Riverside builds off this

previous work and provides dynamic SA through data-driven, real-time visuals of a network's state using mechanisms different from those discussed here.

2.3.2 Industry Tools

Tools like Gephi [79] or Cytoscape [80] are often used to create custom network visualizations but can have additional software or hardware requirements and necessitate the manual creation of the visualizations by end users. Enterprise tools like Splunk are often used to visualize network security events, but the visualizations are an additional add-on that users must create themselves [81]. Arkime [82] and Kibana [83] are open-source tools, similar to Splunk, and are primarily for data aggregation. Arkime focuses on providing packet capture analysis and visualizations versus Kibana's dashboards of network security chart visualizations.

VizAlerts is a "data-driven automation platform" that can be connected to Tableau to build various charts and dashboards for a variety of network alerting [84]. Some would also consider Shodan a geographic cybersecurity visualization tool for IoT devices because its data can be used to create unique and user-specific visualizations [85]. Skydive is another open-source project for real-time analysis, but its centered around Linux operating systems and internal network analysis [86]. GrassMarlin is a network mapping and SA tool, but it is built for ICS and SCADA networks [87].

2.3.3 Cybersecurity Visualization Evaluations

Other research has focused on developing evaluation frameworks for cybersecurity visualizations. Sethi et al. recommended evaluating security visualizations through a cyclic 4-stage

model [88, 89], while Sharafaldin et al. used 7 criterion to quantitatively score network security visualizations through surveying current network security visualization research [90]. Sethi et al. validated their framework with input from visualization and security professionals and base the start of their model off of task goals and data types. Sharafaldin et al. evaluated 26 network security visualizations with their taxonomy and presented their results along with the current challenges in evaluating cybersecurity visualizations, some of which were the lack of validation evaluation and benchmark datasets.

2.4 User-Centered Approaches to Cybersecurity Visualization

While some of the visualization tools mentioned previously incorporated user-driven requirements, there exists a body of work that focuses specifically on incorporating users into the design and evaluation of security visualizations due to the highly technical needs of the cybersecurity field. McKenna et al. stressed the importance of user-centered design methods for cybersecurity visualisations and developed a set of user profiles to identify needs and use cases for when visualization developers don't have direct access to security personnel [14]. Building on this work, McKenna et al. built a cybersecurity dashboard, BubbleNet, using human-in-the-loop development at each stage of their process, conducted a usability study, and deployed their dashboard in an operational environment for further testing [91]. Similarly, Stoll et al. recommended a "persona" approach with a 5 step process to determine cybersecurity visualization requirements and use cases [92], while another study used focus groups and semi-structured interviews with security professionals to apply visual analytics to malware analysis [93] using the "data-users-task analysis" framework [94] for their design process.

I differ from these studies in that I chose to interact with professionals who use visualizations across a variety of network and security roles. I didn't want to rely on generic models that can't account for all possible cybersecurity use cases. Best et al. conducted interviews and focus groups to categorize user groups and the challenges associated with designing security visualizations for those users [20]. To address these challenges, they built a network security visualization mock-up called SEQVIZ by incorporating user recommendations throughout the entire design and development process. Similar to Best et al., Fink et al. conducted an ethnographic study of cybersecurity analysts to understand the challenges they face and provided a set of design principles for building an all-encompassing cybersecurity visualization environment [15]. Erbacher used CTAs along with discussions from various stakeholders to develop Cyber Command Gauge Cluster (CCGC) meant to provide mission impact SA [71].

Arendt et al. [21] used in-situ observations and interviews with security analysts to determine user requirements for building multiple network visualizations with various use cases. Legg used survey data from non-expert users to develop a web-based, cyber-SA dashboard, ePSA, that used a VPN infrastructure to display information about people's personal and home devices [95]. Similar to Riverside, ePSA uses force-directed, node-link diagrams to display its network view and provides some user-customization features to highlight certain activities. Last, a similar study to mine was focused on the need for usable visualizations in the field of system administration, where researchers interviewed system administrators to gather technical visualization features and tie them to domain-level tasks [96]. Franklin et al. used methods akin to mine through group-focused semi-structured interviews with security analysts [3]. Their findings focused on developing tasks from analysts daily workflows and alert triage and analysis tasks, all with the goal of building an alert management visualization tool. I use this prior work as a framework

for how to conduct my research by incorporating user-centered methods to design and evaluate a network security visualization tool.

2.5 Taxonomy of Cybersecurity SA Visualizations

Last, I want to provide a taxonomy of previous work that shows where Riverside differs from other visualization research and tools, similar to Beck et al., who focused on dynamic graph visualization methods [23] instead of specific tools. The taxonomy is shown in Figure 2.1. I do not list all of the visualizations mentioned previously but rather those that focus on the use cases of situation awareness, as well as some of the more common industry tools that provide security-focused visualizations. Some of the tools do not mention specific capabilities directly or to the extent where I felt comfortable to categorize it. The data categorization of "static" versus "dynamic" is meant to show what the tool inherently supports in terms of providing static or dynamic visuals. I color these instances, as well as partial implementations, yellow, to signify that it is something not inherently offered or never clarified. This in and of itself is a finding, as many of these tools which are meant to be "used" do not discuss how an end user would deploy the tool in an operational capacity. Additionally, I acknowledge that this taxonomy could include unforeseen biases, but I present this as a way to visually show where Riverside differs from other work.

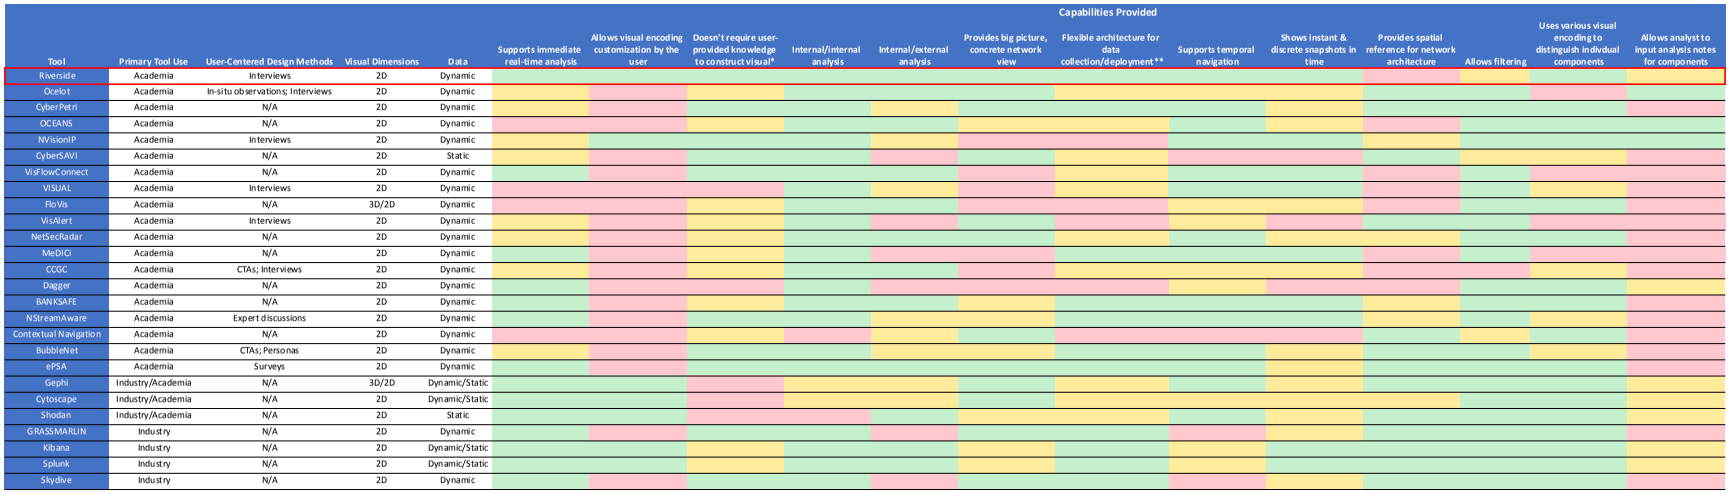

					Capabilities Provided												
Tool	Primary Tool Use	User-Centered Design Methods	Visual Dimensions	Data	Supports immediate real-time analysis	Allows visual encoding customization by the user	Doesn't require user-provided knowledge to construct visual*	Internal/internal analysis	Internal/external analysis	Provides big picture, concrete network view	Flexible architecture for data collection/deployment**	Supports temporal navigation	Shows instant & discrete snapshots in time	Provides spatial reference for network architecture	Allows filtering	Uses various visual encoding to distinguish indivdual components	Allows analyst to input analysis notes for components
Riverside	Academia	Interviews	2D	Dynamic													
Ocelot	Academia	In-situ observations; Interviews	2D	Dynamic													
CyberPetri	Academia	N/A	2D	Dynamic													
OCEANS	Academia	N/A	2D	Dynamic													
NVisionIP	Academia	Interviews	2D	Dynamic													
CyberSAVI	Academia	N/A	2D	Static													
VisFlowConnect	Academia	N/A	2D	Dynamic													
VISUAL	Academia	Interviews	2D	Dynamic													
FloVis	Academia	N/A	3D/2D	Dynamic													
VisAlert	Academia	Interviews	2D	Dynamic													
NetSecRadar	Academia	N/A	2D	Dynamic													
MeDICi	Academia	N/A	2D	Dynamic													
CCGC	Academia	CTAs; Interviews	2D	Dynamic													
Dagger	Academia	N/A	2D	Dynamic													
BANKSAFE	Academia	N/A	2D	Dynamic													
NStreamAware	Academia	Expert discussions	2D	Dynamic													
Contextual Navigation	Academia	N/A	2D	Dynamic													
BubbleNet	Academia	CTAs; Personas	2D	Dynamic													
ePSA	Academia	Surveys	2D	Dynamic													
Gephi	Industry/Academia	N/A	3D/2D	Dynamic/Static													
Cytoscape	Industry/Academia	N/A	2D	Dynamic/Static													
Shodan	Industry/Academia	N/A	2D	Static													
GRASSMARLIN	Industry	N/A	2D	Dynamic													
Kibana	Industry	N/A	2D	Dynamic/Static													
Splunk	Industry	N/A	2D	Dynamic/Static													
Skydive	Industry	N/A	2D	Dynamic													

Offers feature in entirety

Does not offer this feature

Partial feature implementation or not directly mentioned

*This means that user intervention is not required to create the visual or separate distinct components, as long as the initial data collection has been deployed or done.

**This indicates that deployment and a full architecture solution is discussed to be used operationally. For industry tools, this might depend on how it's initially configured.

Figure 2.1: **Taxonomy Comparison of Cybersecurity SA Visualizations.**

Chapter 3: Interviews with Computer Network and Security Professionals

3.1 Overview

The use of visualizations in cybersecurity is not new, including those directly supporting the concept of situation awareness [13]; however, oftentimes these visualization tools are developed without regard to the direct needs or wants of security personnel [14], or the visualization piece is an add-on to a capability that performs a different function all-together, such as the case with tools like Splunk [81]. Furthermore, many of the tools that meet all of these needs are costly and unobtainable outside of a large enterprise environment [97]. With this in mind, I wanted to gain a better understanding of what those exact needs are and how they can be met through the development of a network security visualization tool.

While I recruited participants from a variety of backgrounds, my goals were three-fold: 1) understand the experiences current network or security professionals have about the tools they use, 2) evaluate the level of cyber-SA that participants believe they have of their networks, and 3) gather various perspectives to use towards development of a network security visualization tool. I acknowledge that the needs of a NOC or SOC analyst and network administrator or incident responder vary, but all of these fields use visualization tools to aid in their job tasks. Furthermore, for many small to medium-sized businesses, these roles and responsibilities overlap much more, sometimes to the extent that network and security administration roles fall to the same team or in-

dividual. In fact, some of my participants were the sole IT and security administrator responsible for both network and security operations on their company's network.

3.2 Methods

After receiving approval from the University of Maryland IRB, I conducted 24 audio-recorded, semi-structured interviews with participants who had backgrounds working in a NOC, SOC or experience with network administration over Zoom video conferences. The general flow of interview questions was background information on participant experiences, questions about capabilities they've used, and general questions on good incident response practices and situational awareness in security. The final part of the interview then asked participants to either draw or explain what they need from a visualization tool to help them be effective in their role, which could have involved a layout or specific features. I used Zoom transcription to generate initial transcripts and then used the audio recordings to correct any errors in the transcripts. I anonymized all transcripts by removing any identifying information about participants, such as specific companies or organizations, before importing the transcripts into my coding software. Additionally, all participants were required to sign and return a consent form before their interview began. All of my interview questions are located Appendix A.2.

3.2.1 Participant Recruitment

I recruited participants with at least 1 year of experience working in a NOC, SOC or as a network administrator. I did this primarily through social media posts on forums such as Twitter, LinkedIn, and Reddit, as well as emails through personal and professional networks. My inter-

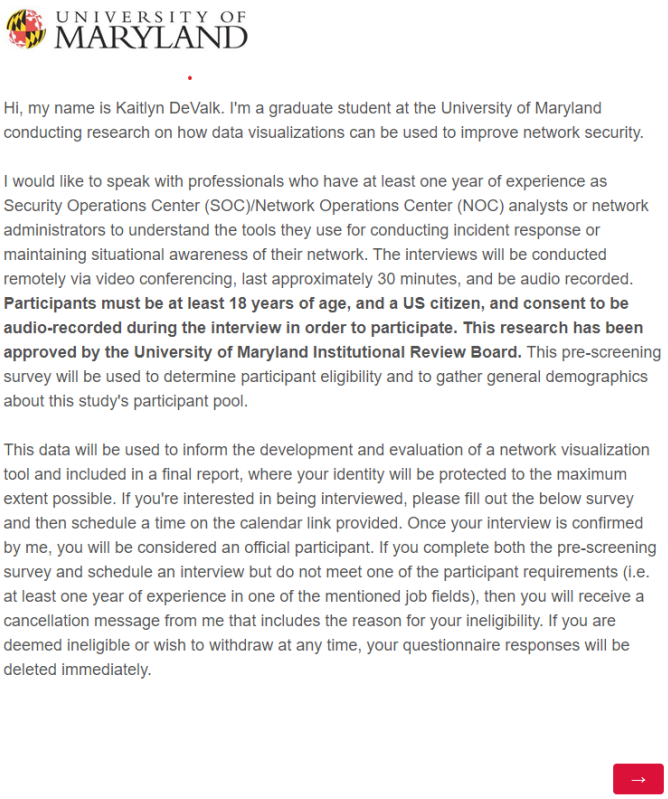

Figure 3.1: **Interview pre-screening survey.** Initial message in the pre-screening survey that had to be read and filled out by participants to schedule interviews.

view recruitment templates are located in Appendix A.1. I used a pre-screening survey through UMD Qualtrics to provide an overview of the study, located in Figure 3.1, which first asked if participants were at least 18 years-old and U.S. citizens. If they were, the survey displayed four demographics questions, which are included in my full set of interview questions in Appendix A.2, and a calendar link to schedule an interview. Upon review of a participant's survey responses, I either reached out for clarification on their experience because of a note left in the submission, confirmed their interview, or removed them from the study. The only disqualifying criteria that I encountered for a participant was lack of experience in a relevant professional field.

I recruited and conducted interviews during May and June 2022 and completed 24 inter-

views during this time frame. Table 3.1 contains basic demographics for all of my participants as well as information about their current professional roles and what area of industry they currently work in. My participants came from a variety of backgrounds with majority of them working directly in the cybersecurity industry in some capacity. Their average experience was about 9 years, and majority of my participants identified as male. Overall, I obtained an incredibly diverse participant pool with varying backgrounds, which I believe aided my end result.

Table 3.1: **Interview Participant Demographics.** I show, for each participant, their gender, years of relevant industry experience ($>$ 1 year), their highest completed level of education (high school, bachelor's, master's, or doctorate), and the current industry they work in.

P#	Gender	Years of Experience	Education	Current Industry
1	Male	1	HS	Cybersecurity
2	Male	8	HS	Information Technology
3	Male	2	MS	Cybersecurity
4	Male	1	HS	Telecommunications
5	Male	5	MS	Finance
6	Male	10	HS	Cybersecurity
7	Female	35	MS	Education
8	Female	5	MS	Cybersecurity
9	Male	15	PhD	Information Technology
10	Male	3	HS	Cybersecurity
11	Male	4	BS	Pulp and Paper Industry
12	Male	6	HS	Information Technology
13	Male	30	BS	Education
14	Male	18	HS	Consulting
15	Male	7	HS	Cybersecurity
16	Male	1	HS	Cybersecurity
17	Male	24	BS	Finance
18	Male	10	BS	Cybersecurity
19	Male	8	MS	Cybersecurity
20	Male	18	MS	Education
21	Female	6	HS	Cybersecurity
22	Female	2	BS	Cybersecurity
23	Female	2	HS	Security Monitoring
24	Male	1	BS	Cybersecurity

3.2.2 Data Analysis

I performed qualitative analysis for my interviews, with some semi-quantitative analysis when looking at how many times a particular visualization feature was coded. My qualitative analysis was done using open coding and thematic analysis, loosely following the framework proposed by Braun and Clarke [98]. I deviated from this framework by finishing my coding and subsequent codebook completely before diving into the thematic analysis.

My qualitative analysis was performed using NVivo [1]. I used open coding and thematic analysis to distill participant responses into 54 codes and 14 corresponding themes. The final codebook can be found in Appendix A.4. I coded 24 "Visualization Features" which were not included in my thematic analysis and are thus not included in my previously mentioned totals. I chose not to use a reliability metric as I was a single researcher, and my analysis was largely qualitative [99]. I wanted to assess the needs that current computer network or security analysts have and what is missing from the capabilities they currently use.

Coding Process

Since I performed every interview and then read through each transcript to correct any errors in the automated transcription, I was very familiar with the interview data; however, I chose to use the process of memoing [100] to keep track of my thoughts throughout the interview and analysis stages to ensure consistent intent behind my codes, as well as a log of how my codebook changed throughout the course of coding my interview data. A small snippet of this can be seen in Figure 3.2. I chose four interviews at random to generate my initial codebook,

[1] https://www.qsrinternational.com/nvivo-qualitative-data-analysis-software

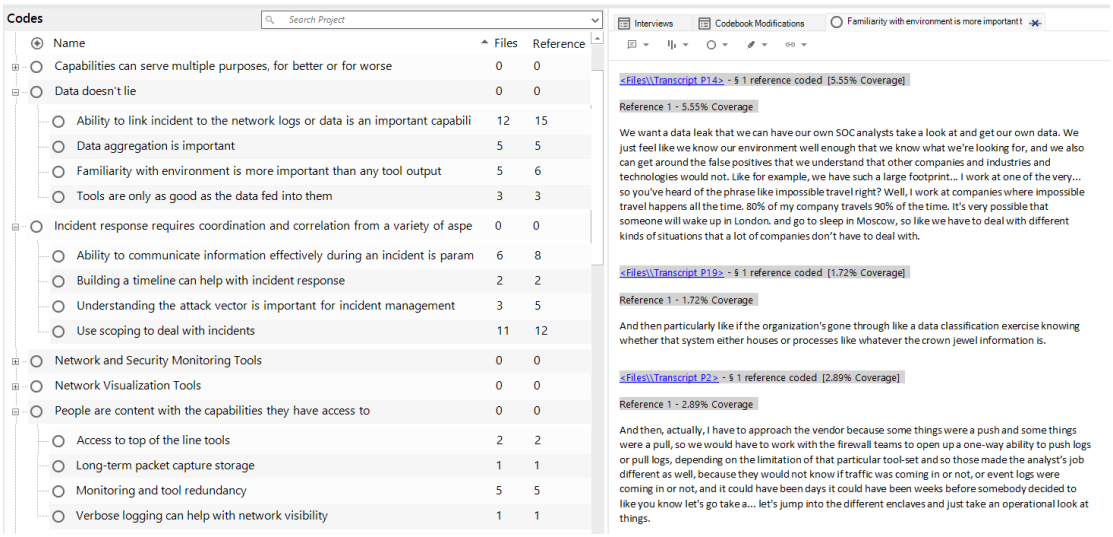

Figure 3.2: **NVivo Codebook Excerpt.** Snippet of codebook with my codes and subsequent themes (left), and quotes from various transcripts for the code "Familiarity with environment is more important than any tool output" (right).

tested it on 2 interviews, and revised the codebook as I saw fit. I then reviewed it with a second researcher who had access to the codebook and interview transcripts to ensure that my coding methodology made sense for my research questions.

I continued this process by taking 2 interviews at a time to apply my codebook. If I felt it needed revisions, I made notes in my memo as I was coding and only added codes that I believe filled necessary gaps in my codebook. I then modified my codes through merging or deletion after I was done with 2 full interviews to ensure a thorough and reliable process where I could review my memo notes from the interviews I had just coded. I did a final review of my codebook, looking through the descriptions and corresponding coded text to ensure accuracy for each code. During this process, I split my larger codes into smaller ones as I felt they weren't specific enough to have meaning. Once this was complete, I once again reviewed the codes with the other researcher as a final check in my coding process. After a bit of discussion, consensus was reached, and I froze my final codebook with 54 codes. I will note that 4 of my interviews

made no changes to my codebook throughout the coding process.

Theme Analysis

Through the use of memoing, I kept track of initial themes that stood out to me during my coding process. Once I finalized my codebook, I wrote out 12 initial themes I'd noted through coding. I then reviewed the themes, combining, removing, or adding new ones as necessary, to ensure that every theme was relevant to my research questions and the codes those themes encompassed. Before completing my thematic analysis, I again reviewed my themes and the codes assigned to each with my fellow researcher.

I finished my analysis with 14 themes. Upon finalizing my themes, I organized the themes into four categories that aimed to provide the necessary background for my thesis and answer to my two overarching research questions. Every code was assigned to a relevant theme which can be seen visually in the four theme maps in Figures 3.3, 3.4, 3.5, and 3.6.

3.3 The Landscape of Network and Security Capabilities

Before discussing my interview results, I want to highlight the vast landscape of network and security tools that were mentioned by interview participants. Over the course of reviewing interview data, I kept a running list of both network and security monitoring tools and network visualization tools that participants specifically mentioned as capabilities they used frequently or were most familiar with. A few of the capabilities mentioned are no longer managed or used operationally, but participants mentioned them to provide context. The tools mentioned, as well as how many times they were mentioned across all participants is located in Table A.1. The

total number of tools mentioned by participants was 54 and was a mix of both commercial and open-source tools.

Additionally, I kept a running list of visualization tools that participants specifically mentioned that they'd used for network visualization capabilities which is located in Table A.2. In some cases, such as with SolarWinds, the tools in the two tables overlap, due to their multi-faceted capabilities. Other tools, like Elastic, can integrate visualization capabilities on top visualization, such as Kibana, where the primary function is providing visualizations for security data. Something to note is that many of my participants couldn't name a specific tool when discussing a particular capability because it was either something custom-built or was something they used occasionally and not part of their daily processes. A few participants refrained from mentioning explicit tools, largely due to their employers, so these lists are only representative of the tools directly mentioned by participants. In contrast to the many network and security monitoring tools discussed by participants, there were only 19 visualization capabilities participants referenced across the corpus of interviews.

3.4 Results

I now present the results of my analysis, which contains my thematic analysis, as well as my qualitative analysis on the 24 visualization features I coded. I did not put the visualization features into themes, as those were coded to be used towards the development of a network security visualization tool.

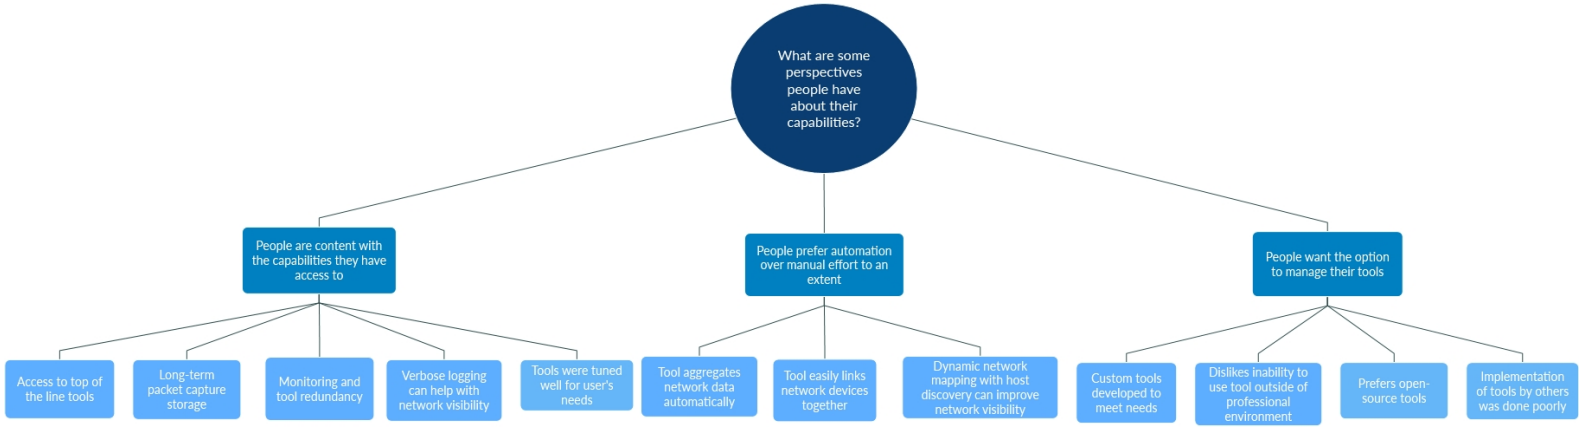

Figure 3.3: **Theme map 1.** What are some perspectives people have about their capabilities? Initial connections are themes, and secondary connections are the codes related to those themes.

3.4.1 Themes

Upon completing my thematic analysis, I realized that my themes shared commonalities, so I organized them into four groups. Each section below contains a theme group which can be seen visually in Figures 3.3, 3.4, 3.5, and 3.6.

What are some perspectives people have about their capabilities?

Since I was directly asking participants about their experiences, opinions emerged naturally, but it was interesting to see the various perspectives participants presented.

People are content with the capabilities they have.

A few participants very directly stated that they were quite content with their capabilities. The reasons behind this sentiment varied, but generally centered around access to "top of the line" tools and the feeling that they had enough overlap and redundancy in their capabilities. P1 stated specific capabilities that made them comfortable with the visibility they have on their network, while P5 felt their company had great redundancy and overlap in capabilities saying *"it's kind of like, you put Legos together and like a little bit overlap, but, eventually everything gets covered."* Other participants, like P12, felt like their company had no blind spots and *"there's really nothing that [they] are wanting."* Other participants, like P17 felt that they had their tools and alerts tuned really well, providing *"great threat awareness"* overall.

People want the option to manage their tools.

When discussing how participants felt about their capabilities, many held the notion that they weren't satisfied with their tools or capabilities because *"they were users of someone else's design"* and couldn't make changes that would enhance their visibility and insights. This particular notion seemed to exist because the analysts, *or users,* couldn't make the changes they wanted on their end and needed a third-party to implement them, which sometimes never happened.

Because of this, P1 and P4 stated that they preferred open-source tools because they can add to it *"on the fly,"* giving them more control over their tools, while P18 stated that their company will develop custom tooling to meet their needs directly. Along this same line, participants like P20 and P22 expressed that some commercial tools are great, but depending on the version, certain components weren't able to be modified or adjusted for their needs.

People prefer automation over manual effort, to an extent.

Despite people wanting to have control over their tools, many of them expressed the desire for an automation component in their capabilities. P3 and P19 both agreed that dynamic host discovery was a common problem they'd faced over their careers and having a tool that could do that would be game-changing. Many participants also said that automated data aggregation was imperative for their operations, and P2 stated that tools aimed at data aggregation *"made the consolidation and just event tagging and correlation of different sets of logs from different technologies together, much easier."*

Participants expressed a likeness for automation particularly when it came to network mapping or visualization capabilities. P12 discussed an automated mapping tool that *"builds an ac-*

tual model of [the] network" and stated that it basically *"does all the work for you,"* which made their job easier. On the other end of the spectrum, P11 said they were *"mostly managing [their] inventory in Excel spreadsheets"* and that they *"don't really have a dynamic inventory."* Furthermore, they expressed the desire for *"something that could build a network map or a node map right of the network using SNMP, or at least you know, even if I have to draw the lines myself. I mean our network is small but with a massive company... that would be so helpful."*

What are some industry truths?

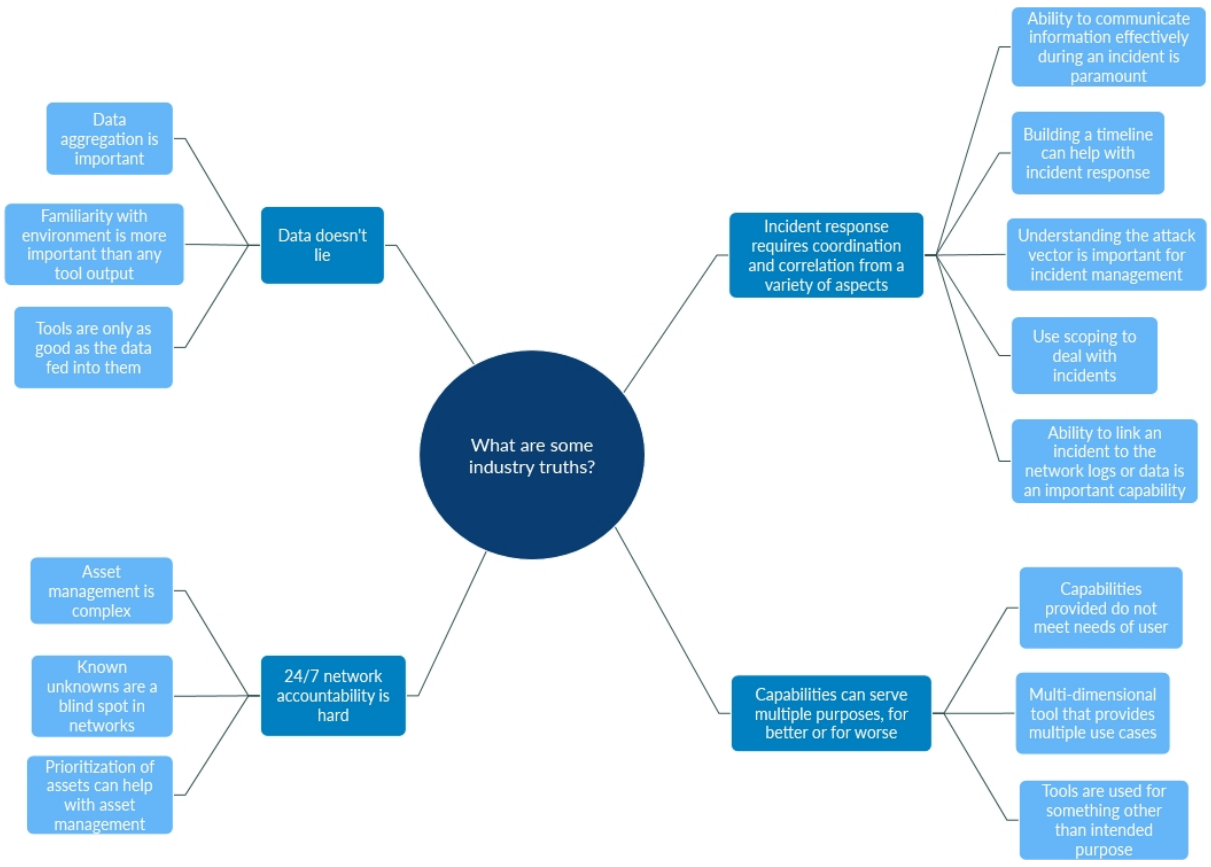

Figure 3.4: **Theme map 2.** What are some industry truths?

Participants repeatedly made statements that there were certain realities when you work in the security or IT field. The themes presented here encapsulate the truths mentioned throughout

interviews.

Data doesn't lie.

Many participants stated that regardless of what tooling you have access to, it's only as good as the data that you collect, and if you aren't familiar with the environment you're operating in, then you can't possibly understand what your data is telling you. P19's stated *"it's kind of the problem with any kind of network monitoring, visualization stuff, it's like, as long as you're feeding it good data then like, it's going to be able to present good data,"* showcasing the importance of maintaining data integrity across your environment and tooling. P14 made a point about the importance of data integrity, as well, stating that *"[I] have a guideline that I tell all of my analysts, and I've used it for a while. People lie, computers can lie, my dog lies, everyone lies. Data doesn't."*

Several participants also felt that having the raw logs or data all in one place was going to be more useful than any tool output, stressing the importance of data aggregation and familiarity with the operational environment. P6 felt it was more important to be familiar with the operational environment instead of one's tools and discussed the importance of *"local and global prevalence.* They stated that *"anything that you can do to combine what you're looking at with some type of global prevalence or local prevalence is very powerful. So one, how common is this thing for this system? That's kind of your local prevalence... or how common is it within this customer or this organization? Global prevalence is really only something that you can get when you kind of monitor, multiple different environments, so how common is this thing across all of my customers, or all the places that I'm looking?"* While participants made several varying points, the bottom

line was that your data should be considered ground truth over what any tool tells you.

Incident response requires coordination and correlation from a variety of aspects.

IR is a complicated task that requires a variety of tasks to be successful, which was shown in the various responses and sentiments participants provided when asked what was *the* most important. Many of my participants said that proper scoping was paramount because as P14 said *"it's making sure that we get the entire story."* P18 echoed this sentiment believing that *"if you focus too micro, you fix the one host, and then you don't realize it's on 10 others, the problem is actually much bigger than you thought,"* while P17 said *"knowing how many endpoints, and like the real number, not the one they write on the whiteboard."* Others, like P13, cared more about being able to track down the data that triggered the alert or incident because *"you get to a certain point that's not necessarily a breach yet until there's some sort of data exfiltration and that all goes back to logs."*

P8 stated that identifying the attack vector was paramount because *"the main goal of incident response is to put in place security controls that will prevent a similar incident from reoccurring after you remediate the current incident,"* while other participants believed that effective communication both up and down the chain was the most important thing during an incident. In that vein, P7 went on to say that having an incident response plan is not enough but that *"you need to test the plan and test the team as well to make sure that everybody knows what their responsibility is in the event that you do have an incident."* This just goes to show the very complex nature of incident response and that people think having that big picture view can really help during a time when the one thing you don't have is time.

24-by-7 network accountability is hard.

Several participants felt that even with the best tools or analysts, certain networks and systems are just complex, making good security hard. Many participants held the notion that there were just things on their network that they didn't know about and weren't going to know about until there was an incident or event that would cause them to "find" that host. P19 described it *"like if the SOC works an investigation and finds out about some weird asset, like a vulnerability scanner,"* that was previously unaccounted for on asset lists or network diagrams, it could be added for the next time, but that "known unknowns" are just a reality in the world of network security.

P8 felt that prioritization of assets would help *"even if you don't have an inventory of all several 100,000 workstations, servers, if you can just get your 50 most prized systems"* than it's better than nothing, especially for larger networks, while P19 felt that *"relying on a tool or person for complete network topology or asset management is not going to work all of the time, especially for larger networks."* P6 summed it up rather succinctly with the sentiment that *"asset management is always crazy."*

Capabilities can serve multiple purposes, for better or for worse.

A common sentiment that seemed to be split amongst positive and negative reactions was that participants felt that they had capabilities that provided multiple functions within their environments or processes, but these capabilities might not be what they really needed. P5 stated that they preferred one of the tools in their environment because it was used as the back-end in other capabilities they possessed, making it a multi-dimensional tool that supported a variety of

functions for an analyst, and P20 said they leverage the unintended functionality of certain capabilities to their benefit, particularly for network perimeter insights. While these participants had positive experiences with tools that they felt provided useful yet unintended functionality, some participants expressed frustration.

P3 felt they lacked the ability to see what was happening on a host at a given point in time because they didn't have a capability that gave them ground truth on every segment of their network. P2 mentioned that they *"would actually use the SIEMs, not for security purposes, but for operations purposes, because they may have had capabilities that we were lacking in our current environments"* and that *"all of these tools [they'd] worked with were never implemented in the way they were designed to be implemented... jerry-rigged, if you will."* P9 and P21 stated that they were provided capabilities that didn't exactly meet their needs, whether that be because they worked in a large organization that mandated certain tools for multiple teams, within a small team with a limited security or IT budget, or with customer-provided data that just didn't supply enough visibility across operational environments. P16 felt that their network visibility was limited because of the capabilities provided and said it was like *"looking at an entire Where's Waldo page through a pin hole. It takes so long to tie things down because we can't see big picture."* Participants like this felt handcuffed at times when attempting to perform their jobs because they had just enough insight to "get it done" but not enough to ever feel completely comfortable operating within their own environment.

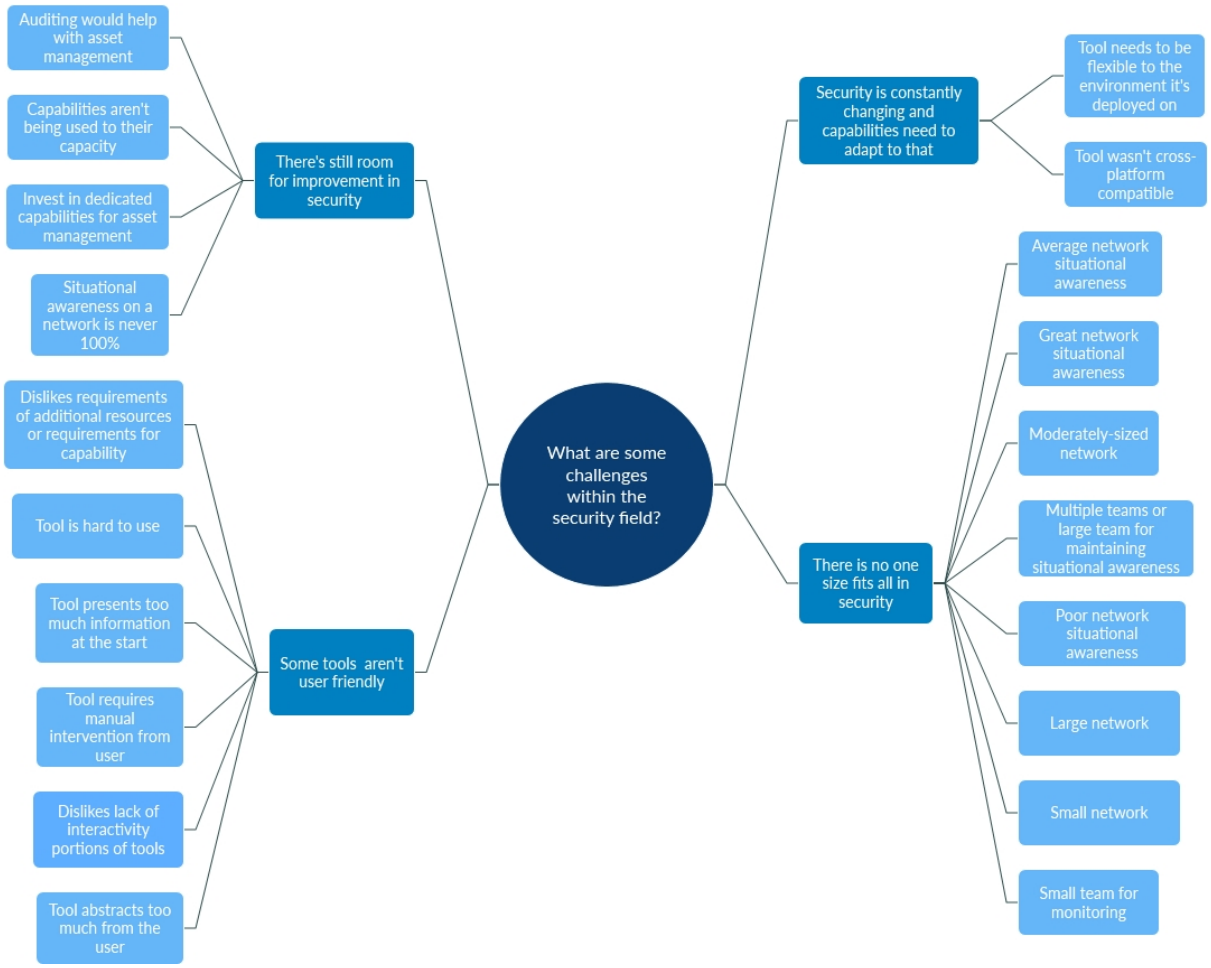

Figure 3.5: **Theme map 3.** What are some challenges within security field?

What are some challenges within the security field?

This next group of themes focuses on the notion that personnel, tools, and the security industry still fall short and discusses specific aspects that participants mentioned during interviews.

Some tools aren't user friendly.

Participants were quick to communicate frustrations with network and security monitoring capabilities that they felt just weren't easy to use, whether that be because they required additional user input or resources to function, causing more stress on the user and network. Contrary to the

theme discussed earlier, many participants had complaints centering around the frustrations for having to manually intervene when attempting to get useful output from a tool. P13 didn't possess an automated way to create logical and accurate network diagrams and stated that *"it's a manual process to create the links"* between all of their assets. P1 had complaints about a tool that they felt was *"clunky and slow, and [they thought] it could be replaced by much better tooling."* Others felt that certain tools were a hassle because they relied on additional hardware or software requirements that caused performance issues, such as P20 who couldn't capture more data with a tool because it would decrease the tool's performance, making it unusable.

P21 had an interesting comment about visualization tools specifically, in that they *"stay away from visualization tools, because sometimes [I'm] intimidated by it, and it feels like a time sink"* for them to spend time getting proficient at it, and even while some are better than others, most are not intuitive. P17 echoed frustration with security visualization tools because *"the biggest mistake... and where all of them go wrong, is they're just too noisy right out of the gate... there's just a ridiculous amount of data, and I think as people, we start to protect ourselves and shut down when confronted with too much data."*

There is no "one-size fits all" in security.

The size of a security team and the situational awareness provided by their capabilities vary drastically from one organization to the next, and with this, the needs of a cybersecurity company can differ from those of a finance corporation. For example, P10 and P11 worked on networks with about 200 to 300 endpoints whereas P23 cited around 32,000 endpoints for their respective network scope. Additionally, participants, like P24, had very large operational environments but

stated that they only have about 10 analysts that staff their 24/7 operations center, whereas P16 stated that they worked at a corporation where multiple teams are monitoring their network at a time. This isn't to discredit the work being done but just that the needs of one company are vastly different from another and capabilities that work in one environment may not be suited to work in another.

Additionally, participants had varying reasons for why they felt confident or not in terms of their situational awareness. Participants, like P21, work with customers, so they felt this limited their visibility and accountability of the networks they operate on because *"it's what [they] give us."* Others like P4 felt that they had great visibility because of their tool stack, however many participants felt that their network visibility and accountability was average with a common answer being that their asset management *"as accurate and up to date as possible"* as stated by P10. This further showcases a security capabilities gap whether that be from tooling or personnel, which isn't something that an analyst or team can control at times.

There's still room for improvement in security.

Many participants, even those that were confident in their visibility or satisfied with their capabilities, felt that there were still areas that could be improved upon when it comes to network visibility and overall SA within their environments. P11 was content with the insights their tools provided but wasn't *"really satisfied with where [they're] at with using those tools,"* which was similar to P16 who felt that their team didn't know how to use current capabilities to the best of their ability. Others like P2 and P4 felt that continuous auditing and verification of network diagrams or asset inventory would have greatly improved their visibility. P10 believed that it

"boils down to communication, whether it's internal communication about who's using what resources or communication with our clients as to what resources they're using, what products." At the end of the day, P7 put it perfectly when they stated that *"nothing is ever 100%."*

Security is constantly changing, and capabilities need to adapt to that.

While some participants such as P17 felt that over time they were able to adapt their capabilities to their needs, others, like P2, felt that their tools didn't withstand the test of time as their environment grew and evolved. P3 specifically felt that current security visualizations weren't adaptable to the environment they operated on since *"50,000 assets is kind of hard to visualize. You don't want to visualize all 50,000, but we also don't want to basically only visualize 5."* P10 and P23 expressed the importance of "customer-focused" tooling and that sometimes vendors aren't amenable to changing a tool to fit the needs of a company, which can impact their visibility. P18 made a good point that *"it's that cybersecurity thing... attackers figure out something new, and you've got to pivot and build something else,"* which makes designing adaptive capabilities a challenge.

How can information visualizations aid security analysts?

Since I knew that I wanted to design a visualization tool, I focused a portion of the interview on network and security visualization tools. With this, participants described their experiences with such tools and expressed where they felt these tools fail and succeed.

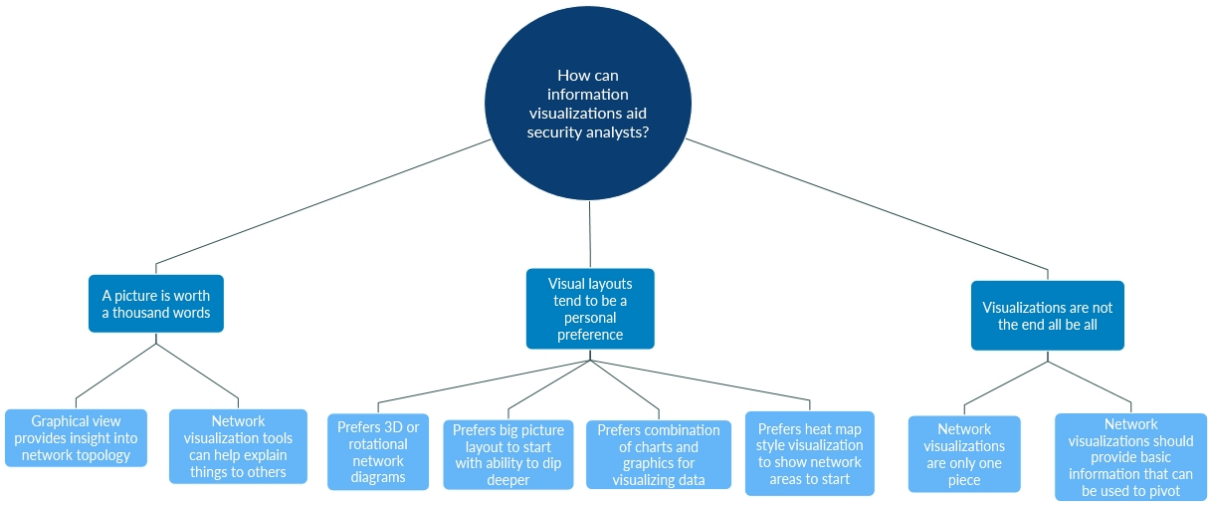

Figure 3.6: **Theme map 4.** How can information visualizations aid security analysts?

A picture is worth a thousand words.

When discussing if participants had used a visualization tool to assist in their professional lives, many of them expressed that having that visual component can help translate a network from just IP addresses to an actual network topology. P1 believed that having that having a graphical perspective provides insight to *"not just network topology but also how segments of the network are used,"* and P22 felt they they can give a *"whole holistic picture versus like sometimes you get a little rabbit-holed or pigeon-holed when you're like down in the weeds."*

P21 felt that *"the best value out of [network visualizations], is trying to translate what we're doing and what we're seeing to someone who doesn't sit in that position."* Additionally, a few of my participants were supervisors at this point in their careers, so they saw visualizations as a way to communicate with upper leadership who may not be a technical by trade but needs to understand the impact of a network incident. P9 said it's like *"you've got people that will look at you like this is Greek, and you've got others that completely speak the language,"* and they see visuals as a way to bridge the communication barrier.

Visual layouts tend to be a personal preference.

People had differing perspectives on what layout they preferred for visualizing a network which varied from heat maps, node-link maps, visuals combined with graphs, or even 3D visualizations. Many participants described their preferred layout as the "big picture" with the ability to zoom in on specific hosts or sites to get more detailed information. P18 preferred the style of heat maps where they would want to see zones of traffic and then upon zooming in to a particular area be able to see specific hosts.

While participants had varying layout preferences, a common theme across all of them was they didn't know what the best layout was for everyone, but the ones described were what they knew and felt comfortable with, such as P12. They went on to say that they didn't know if their preferred hierarchical layout was the best, but they *"wouldn't be able to use a network visualization tool that didn't have those features."* This supports the notion that visualizations need to flexible for the user while also maintaining usability in the features they provide.

Visualizations are not the end all be all.

Some of my participants expressed that, for them, the visualization was only one piece of the puzzle. P14 said network visualizations are *"not the defacto standard because of how much traffic and everything [we] have going on."* P15 felt that visualizations were just extra capabilities, and they relied more on case management tools with automated alerting. These participants felt incorporation of incident management or automated alerting into the visualization pipeline would prove more worthwhile for their roles.

Other participants expressed that visualizations should be able to provide high level infor-

mation that allows them to then pivot to the host level or other tools. P6 provided a great example where *"there might be 50 gigs of traffic between this node and this node"* and being able to visualize that quickly can provide necessary insights when an analysts needs it.

3.4.2 Visualization Features

As noted in my codebook, I coded 24 features that participants mentioned would be helpful in a network visualization tool mapped to low-level actions that can be seen in Table 3.2. I chose to code features generically since this was going to be used toward the development of a new tool versus modifying something that already existed. I coded these features through direct mention of something by a participant, like "filtering," or through an example use case provided by a participant. Descriptions of each mapped task means can be seen in my codebook in Appendix A.4. I used the "Count" column in the codebook, which tracked coded features by participant, so as not to double count a particular feature if a participant mentioned it multiple times. These "counts" should not be seen as statistically significant for a given feature but rather a method that I used to organize and prioritize components that I included in my tool.

Table 3.2: **Visualization Features.** This shows all of the visualization features I coded when discussing network visualization. The second column also shows how many times a feature was mentioned across all of my interviews.

Feature	Number of Times Mentioned
Ability to drill-down for more detailed data	11
Ability to have different views or dashboards for the visualization	10
Ability to visualize network communication between hosts automatically	9
Ability to filter data and save these filters if needed	8
Ability to show network segmentation (visually)	7
Ability to classify or label network entities based on historical data	5

Continued on next page

Table 3.2 – continued from previous page

Feature	Number of Times Mentioned
Ability to add criticality tag to hosts	5
Ability to correlate incident or case management alerts to traffic	4
Ability to label nodes or hosts	4
Ability to have different symbols for different types of hosts	3
Ability to scale or move visualization	3
Ability to see shared resources for an application or host	3
Ability to share layout or display with other users	3
Ability to show amount of network communication visually	3
Ability to use colors to convey some event to user	3
Ability to have multiple environments or profiles	2
Ability to highlight parts of the graph and show stats	2
Ability to show network from a geographic perspective	2
Ability to show snapshots in time of the network state	2
Ability to connect to other APIs or resources to correlate data	2
Ability to see metadata for hosts or machines that are unaccounted for	1
Ability to show general counts of interest to a user	1
Ability to showcase common anomalous traffic to guide users to investigate	1
Ability to collapse or pop-out menus versus static option	1

Overall, the biggest gripe participants had with visualizations they've used was how cluttered or overwhelming many of them are upon start. Consequently, one of the most common features that participants desired was the ability to have basic information presented up front and then be able to get more information about a component at their discretion, echoing the mantra of "overview first, zoom and filter, then details on demand" [101]. This was largely because participants felt visualizations could present too much information up front and lead to "burnout" for analysts. Participants also desired the ability to have multiple views or dashboards since people

process information differently. P22 said sometimes they add charts to their dashboards *"so that it's visually entertaining for [their] eyes, so that [they] don't get burnt out when looking at [the] graphs."*

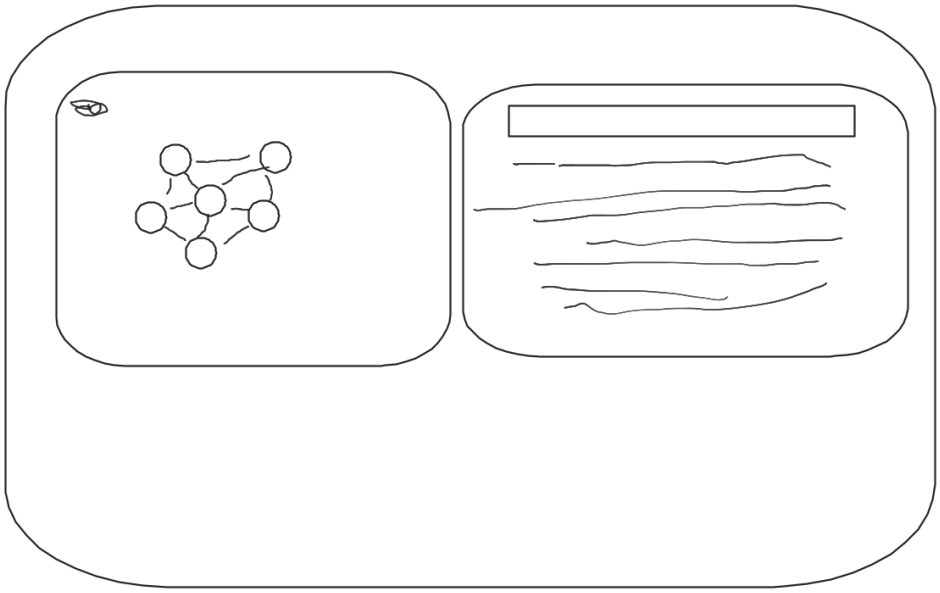

Figure 3.7: **P1 Sketch.** Depiction of network visualization tool layout by P1, showing a portion of the network as nodes connected with edges to represent communication (left) with detailed information about that portion of the network (right).

Another feature mentioned by participants criticality tags for hosts, and P8 thought this was particularly important for either new employees since they're *"probably not going to know what IP address belongs to a certain server, and what applications are residing on each server, so having that ability to discern that this is a system that holds PII or PHI data and having that, very clearly marked out on the diagrams, is really helpful in my opinion."* A modern feature that participants felt was increasingly important with more complex networks was being able to show network segmentation through the visualization, whether that be done automatically or with a tagging ability provided to the user. One of the other common but more complicated features

was the ability to filter traffic whether that was to just tune out "the noise" or because their jobs require them to look for very specific indicators, such as P24 who feels that having the ability to perform *"log querying and a log search"* is key in threat hunting.

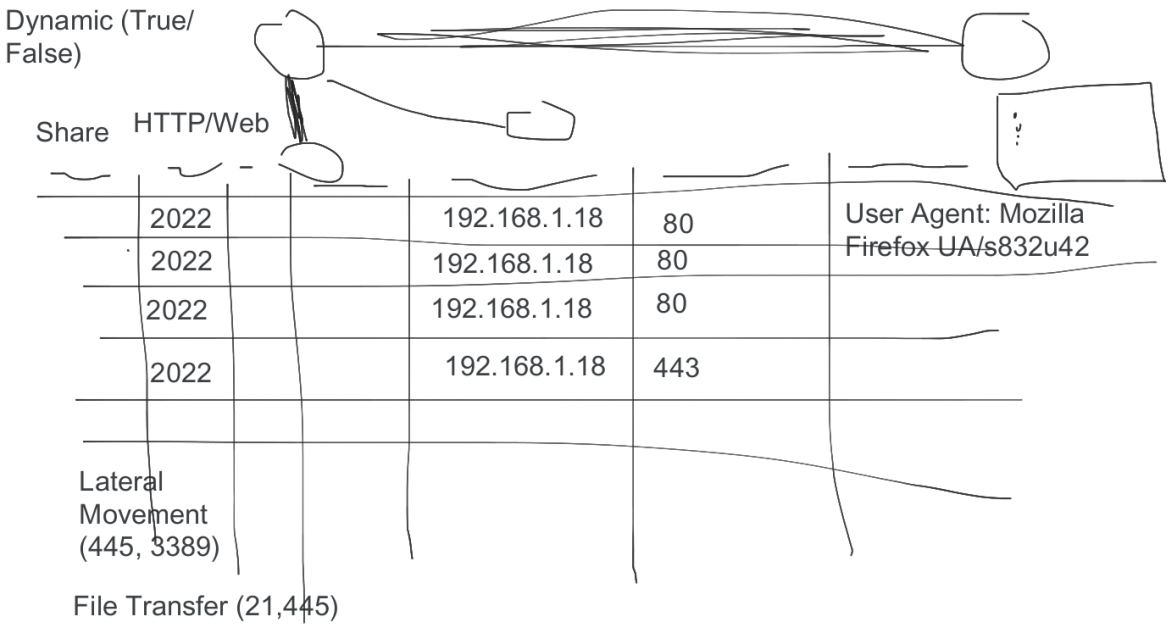

Figure 3.8: **P6 Sketch.** Depiction of network visualization tool layout by P6 that use node-link diagrams (top) for the network visualization along with a chart of network traffic (bottom).

As mentioned before, I gave participants to the option to draw when I asked them the final interview question. For participants that chose to exercise this option, it provided very helpful depictions of what they envisioned for a network visualization tool. P1 preferred a node-link visualization but liked the idea of a feature that allowed them to choose sections of the graph to show more information about those hosts or traffic as seen in Figure 3.7. P6 preferred having multiple views or dashboards that displayed the network both from a visual and chart perspective as seen in Figure 3.8. P6 also liked the idea of visually portraying the amount of traffic by making thicker links between nodes that represent hosts, as well as the ability to turn dynamic updates on or off. While I don't believe that one tool can answer all of the problems or desires presented

by my participants, I do think that the concept of flexible situation awareness tools can provide necessary insights. With this, I started to build out my network security visualization tool using the data collected and analyzed from my interviews.

3.4.3 Limitations

Since I was recording participants and asking about tools they've used, some participants declined to answer certain questions or could not provide clarification because of their professional roles. This primarily affected their ability to name specific tools when asked and did not largely impact their responses for the other questions. I also took care in logically ordering my interview questions from broad to narrow scope when discussing the types of capabilities participants use and the challenges that come with them, but I didn't use methods such as a pilot study to formally test my questions. I did review my questions with a second researcher, but my questions could have contained potential biases through word choices or phrasing.

I acknowledge that using a single coder is not the norm, but as mentioned above, I was the primary researcher where my goal was to uncover themes present in my data. Second, I chose not to use a reliability metric and instead used memoing with review by a second researcher at major points in my analysis. My semi-quantitative analysis for specific visualization features is the main area which could have produced unintentional coder bias. I attempted to mitigate this bias as much as possible through the methods mentioned previously. I also did not use the semi-quantitative portion of my analysis to make any statistical claims. Additionally, my participant pool is not representative of all backgrounds in the technology industry, but it does showcase a wide range of experience, education, and professional backgrounds.

Chapter 4: The Riverside Visualization System

4.1 Overview

There are several challenges that make cybersecurity visualizations particularly difficult, rooted in the myriad of data sources and quality in a given environment [20]. While I only utilize network flow data for Riverside, I incorporate the concept of streaming data [13] by providing analysts real-time updates with the option to traverse previous time by combining two classes of visualizations: node-link diagrams and timelines. The interview portion of my research was used to inform the design of a network security visualization system. Riverside is a web-based cybersecurity visualization system where animated node-link diagrams are used to show dynamic traffic flow over time. An overview of the Riverside visualization can be seen in Figure 4.1, where cyan-colored "agent" nodes represent internal network nodes and gray-colored "remote" nodes represent hosts communicating with the internal network. Riverside's layout is force-directed, as nodes will not be displayed on top of one another, but they can be dragged or moved by a user to change the layout as they see fit. Additionally, when a user hovers over a link or node, they are provided more details about that component. Links, or edges, between nodes are created when network communication occurs between two hosts and is categorized as "to," "from," or "bidirectional" traffic. Riverside uses a time-to-time mapping for displaying visualization components based on the timeline component, allowing users to watch their network state in real-time

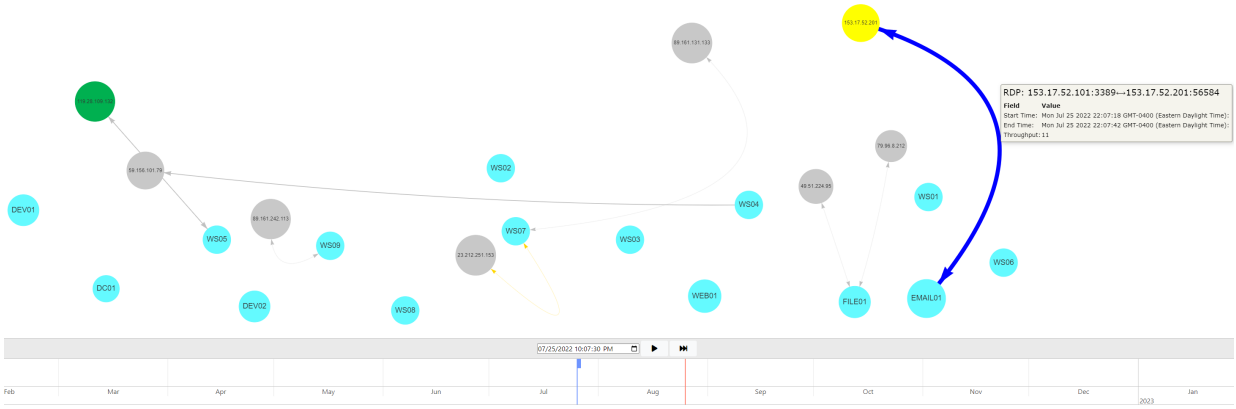

Figure 4.1: **Overview of Riverside Visualization.** Snapshot of network flow showing communication by "remote" hosts (green, yellow, gray) and internal hosts (agents; cyan nodes). Node colors can be changed by a user, like the green and yellow-colored nodes. Visualization components can be hovered over to show more information about a node or segment of network communication, like the RDP traffic displayed. The timeline at the bottom shows both real-time (red cursor) and user-specified time (blue cursor), allowing a user to dynamically navigate their network in time.

or navigate in previous time.

I use animation to visually depict network communication over time. Single view visualizations have been shown to provide faster analysis for analysts [102], while animated visualizations are more accurate for adjacent time and local pattern analysis [103]. While Riverside provides a singular, big picture view of the network, it is best used to compare network topology snapshots throughout time and allow analysts to correlate network communication with potential IOCs. Furthermore, Boyandin et al. stated that animated visualizations need to provide both a play button and slider functionality, giving users multiple options for how they interact with the animation [103]. I incorporate both in Riverside, allowing a user to pause, play, fast-forward, and slide through time as shown in the timeline at the bottom of Figure 4.1. The timeline can also be scaled to allow for finer-grained navigation when dragging the cursor, or the user can pause the visualization and input a time on the box just above the timeline display, but I do not superimpose the animations on to the timeline directly as seen in some hybrid visualizations [23].

4.2 Riverside Use Cases

One of the most important pieces of developing a new capability is identifying its use cases, so the developer can determine what is needed to allow users to be "successful." I used the low-level actions coded from interviews to infer higher level tasks, and since the goal of Riverside was to provide big picture cyber-SA over time using network flow data, it best fit into the SA and correlation analysis tasks [89]. With this in mind, I wrote out some of the initial use cases for Riverside. The examples below are not an exhaustive list but represent both the inspiration for the development of Riverside as well as discussions from interview participants:

1. On-site analyst that needs a "big picture" view of their network to provide situational awareness.

2. Security analyst responding to an incident that needs to determine what portions of the network are potentially impacted through correlation of alerts.

3. Asset management capability with the need to build out an accurate network map and asset inventory.

With these use cases in mind, I lay out the design and development process of the Riverside system.

4.3 Design and Development Process

Riverside's visualization provides a single view that is changed over time. This provides user a simple yet vast interface I use animated node-link diagrams that update in real-time to

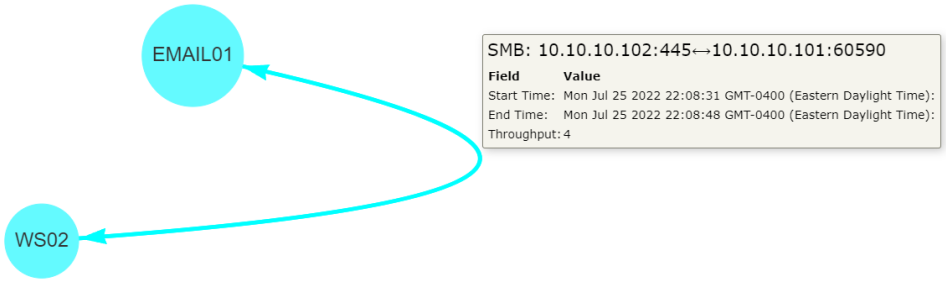

Figure 4.2: **Snapshot in time of Riverside visualization.** This shows two agent nodes (EMAIL01, WS02) communicating over SMB, showing bidirectional traffic on the Riverside visualization and in the hover box.

address the online problem [23] because those are the most explicit representation of a network topology and depict information in a way that security analysts can immediately identify visually, as seen with the sketches interview participants drew in Figures 3.7 and 3.8. While node-link diagrams are common among modern network visualizations, they have proven to be an effective visualization technique to depict dynamic network graphs [23]. Edges will show a single arrowhead on an edge to represent traffic "to" or "from" an agent node, and a two-sided arrow if the traffic is bidirectional, meaning both hosts are communicating with each other as shown in Figure 4.2.

Additionally, I address the transition problem and and aspects of change blindness [104] through animated, temporal navigation [23] by using throughput to change edge thickness based on the amount of traffic and node transparency to show hosts that are communicated with more frequently so that analysts are drawn to "new" or less frequent network communication. Other methods in animated temporal navigation have used staged transitions [105] or graph morphing to preserve mental mappings during transitions [106], but these methods have had varying success in practice. All of Riverside's animations are performed with a time-to-time mapping to assist with temporal navigation, while also providing real-time updates to combat the online problem, giving

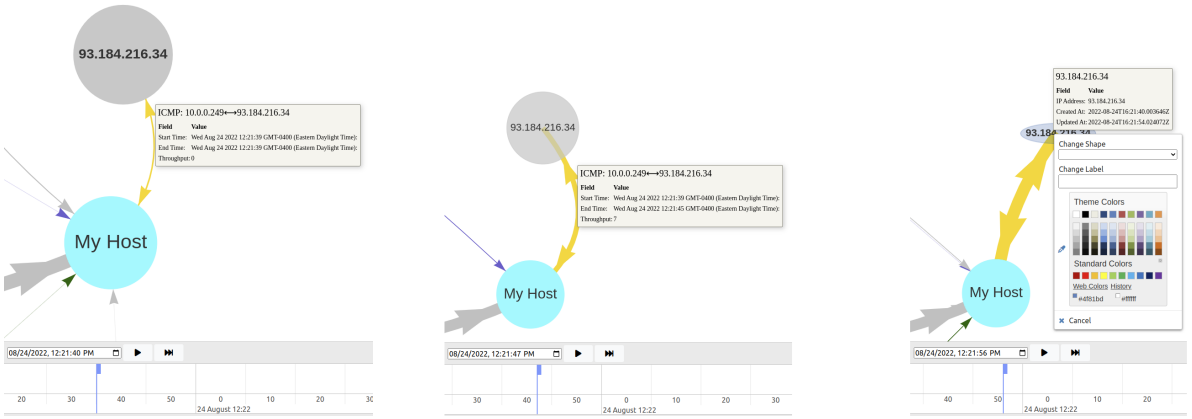

(a) The initial connection shows a yellow-colored, bidirectional edge to represent ICMP traffic.

(b) Edge widths increase as the throughput of network traffic increases.

(c) A node's shape, label, and color by right-clicking it. Node transparency also increased.

Figure 4.3: **Navigation in time with Riverside.** This shows a zoomed in snapshot in time of a "ping" from my host to a remote host. The initial connection shows the start time and end time as the same with a yellow-colored edge to represent ICMP traffic (a). As the connection continues, the edge gets thicker as the throughput of network traffic increases (b). As a remote host is communicated with it's transparency will increase over time (b-c). I can also change a node's shape, label, and color to differentiate it from others on the visualization (c).

an analyst immediate insights into their network. I allow users to alter the visual encodings of nodes through various visual channels, such as shapes, hues (color), and labels to create a special-purpose layout [23] enabling them to track or ignore specific components through time. Edges are colored automatically based on their network traffic protocol, defaulting to grey edges for traffic that cannot be categorized beyond the transport layer. Figure 4.3 contains adjacent snapshots in time that showcases the design components aimed at combating change blindness [104] when using Riverside.

4.3.1 Interview Feature Implementation

The features that I directly implemented from my interview data are listed in Table 4.1, along with the Riverside component that contains that feature.

Table 4.1: **Riverside Features Implemented.**

Feature Implemented	Visualization Component
Ability to drill down for more detailed data	Box appears with metadata when hovering over nodes or edges to get details-on-demand
Ability to visualize network communication between hosts automatically	Automatically connect nodes with edges if network communication is occurring at the time specified
Ability to label nodes (i.e. hosts) on graph	Node labels can be changed by right-clicking one or multiple nodes
Ability to have different symbols for different types of hosts	Node shapes can be changed by right-clicking one or multiple nodes to have different representations
Ability to scale or move visualization	Infinite canvas allows for zooming and dragging to change what is in the user's view
Ability to show amount of network communication visually	Edge width changes based on network traffic throughput, or amount of network traffic
Ability to use colors to convey some event to user	Node colors can be changed by right-clicking one or multiple nodes; nodes are automatically colored based on remote versus internal
Ability to show snapshots in time of the network state	Use of navigable timeline allows user to see real-time and traverse back in time to see how the network state changes
Ability to see metadata for hosts or machines that are unaccounted for	Metadata is pulled for remote hosts and can be used to identify internal nodes that do not have agents deployed
Ability to collapse or pop-out menus versus static option	Uses pop-up window and collapsible menu to let users register and login

The term "drill down" is an information visualization canonical term often referring to obtaining finer-grained details about sub-components in hierarchical visualizations. Here, I use the feature of "drill-down" to refer to the interaction intent of obtaining a "detail-on-demand" view [107] when hovering over components. I was not yet able to implement every feature requested by interview participants, but I was able to include many of the user-specified customization features, which was in line with my thematic analysis results and user-centered design. One feature I would like to highlight is the addition of multiple views or dashboards, which was the second-highest requested feature by participants. While this can be helpful to showcase

the data in various ways, research has traditionally shown that multiple views with multiple attributes in dynamic visualizations can lead to slower and less accurate decisions from users [102]. Some more recent findings have shown that switching between views, vice having multiple views, could actually lead to more positive outcomes, but more research is needed to confirm these results [103]. With this in mind, I felt that further testing and analysis was needed before adding additional views to Riverside's frontend.

4.3.2 Theme Implementation

In addition to the visualization features I directly coded during my interviews, I noticed that my thematic analysis provided important context to as to what people desire from their capabilities: they desire automation but still want to control what a tool shows them; they want a capability that can aggregate data to determine ground truth; and they want something that can be scaled to fit their needs. I used these as overarching guides for initial design and architecture of Riverside.

I worked to address some of the challenges participants faced with their capabilities, particularly that the tool can be used in varying environments and is adaptive to a user's needs. I architected Riverside with native cross-platform capability in mind to ensure compatibility across different operating systems and environments. Additionally, my frontend is web-based, so that it can be accessed from anywhere, allowing users the flexibility in how and where they deploy it. Participants also mentioned certain challenges they had made tools "hard to use," but it can be difficult to capture what specific components make these tools particularly tough to operate. I try to address these challenges by incorporating user-customization components for Riverside's lay-

out and maintaining a simple interface so as to not overwhelm users. The interview participants also came from a variety of industries and felt that some of their capabilities were "one size fits all" solutions which made the tool less effective in their environment. For this reason, I chose to use a client-server model where users can decide what parts of their environment they wish to visualize, whether that be 10 hosts or 1,000 hosts. My utility was purpose-built to encompass themes across network security capabilities and incorporate features requested from interview participants, while maintaining a clean and user-friendly interface.

Since participants felt that visualizations could help paint a picture when explaining their actions or a network incident, I wanted to make sure that the visual component was simple and told a story. Visual layout preferences also seemed to differ across participants, so while I chose node-link diagrams, participants can change the layout of the nodes on the canvas as well as the node characteristics to make the visualization their own. Last, I want to highlight the themes "Visualizations are not the end all be all" and "Data doesn't lie." While visualizations can provide multiple use cases, including those mentioned above, they are not the only tool that security analysts use and that likely will not change. Riverside was developed with the goal of aiding the cyber-SA security analysts have of their environments and augmenting current network security capabilities.

4.4 System Architecture

An overview of Riverside's system architecture is located in Figure 4.4. Riverside uses a client-server model with agents that send network traffic to the server as they receive it using RPCs with Golang protocol buffers for structuring and serializing the data. The server then

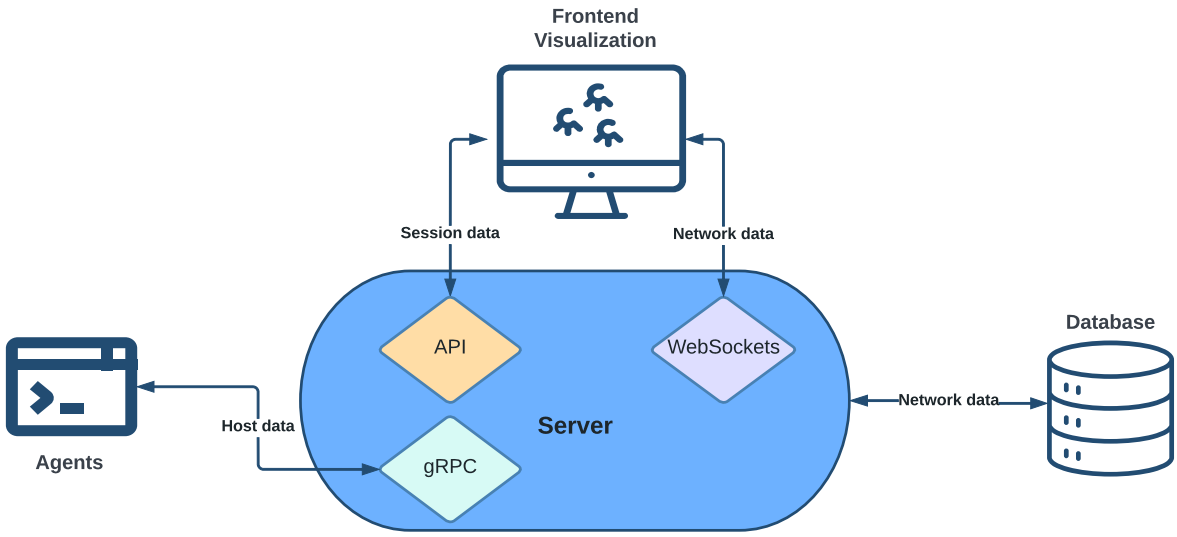

Figure 4.4: **Riverside System Architecture.** Agents are installed on internal hosts that send network flow data to the server as they see it. The server stores the data in a separate database. The frontend visualization uses Websockets to communicate with the server to ensure real-time updates. Session management is handled using an API.

batches the data in two second windows and uses WebSockets [108] to send data in real-time to the frontend visualization. Session management is handled using an HTTP web framework with an API on the backend.

4.4.1 Client-Server Model

Riverside's backend and frontend should be viewed as separate components. While, the frontend communicates with the server to gather data, the functionality of each is distinctly separate. [1]

[1]The backend source code and stand-alone agent and server binaries are located on the Riverside Github: `https://github.com/artemis19/riverside`.

Agents and Server

I use stand-alone agent and server binaries developed in the Golang programming language [109] for the underlying Riverside architecture. Every internal host is monitored using an installed agent that listens on its available network interfaces and sends traffic to the server as the agent receives it. [2] This is done using a client-side streaming gRPC with defined protocol buffers [110] for handling what data is sent and received. The server is responsible for storing the necessary data in the database as it receives data from agents.

Frontend

I used the Javascript library vis.js [111] for the frontend visualization layout. The frontend does not directly communicate with the database to limit the number of concurrent lookups and preserve overall performance. Data is batched and sent to the frontend from the server using WebSockets with the Gorilla WebSocket package [112]. I use the Gin HTTP framework [113] for session management on the frontend. Users will see a prompt like the one shown in Figure 4.5 when registering or logging in. They need to enter a username, password, and server IP address and port. This allows users to manage multiple environments through different Riverside servers.

4.4.2 Database Model

I use GORM [114], a Golang ORM library, to handle database entries and lookups. I used five different tables to store information for Riverside: hosts, network interfaces, users, remote hosts, and network flow. The hosts table was used to store information for agent-installed hosts

[2] Agents require root or administrator privileges to collect host network data.

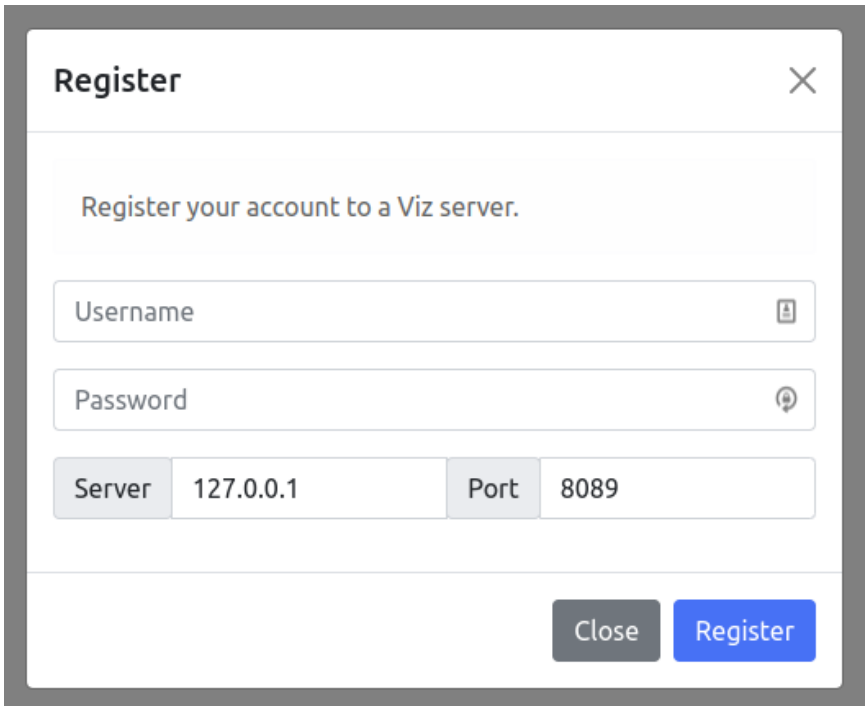

Figure 4.5: **Riverside registration.** This pop-up will appear when a user chooses to either register or login. A username and password must be provided as well as a server to connect to.

on the internal network and contains their operating system, architecture, a unique machine ID, and a link to the network interfaces table. The network interfaces table was used to store all of the interfaces for a given host based on its unique machine ID. The users table contains session information, such as usernames and password hashes. The remote hosts table is used to store every remote host that an agent node communicates with and contains an unique ID, IP address, creation time, and last updated time.

The network flow table contains the most important information for the Riverside system. Each piece of network flow contains two host IDs, depending on whether it's communication internally or between an internal and remote host. In addition to the host metadata, source IP address and port, destination IP address and port, protocol, and start and end times are also stored for each piece of flow data. As mentioned previously, network flow data is batched in two second

intervals when being sent from the server to the frontend visualization. This interval was chosen because it proved to be the most efficient timeframe for determining the direction of the traffic. By comparing each new network communication and checking to see if it is the same as previous traffic, Riverside determines which directional arrow to display on the visualization as well as the throughput for network communication.

The first operational version of Riverside was finalized in August of 2022 and evaluated by real users through a task-oriented usability study.

Chapter 5: Riverside Evaluation & Results

5.1 Overview

A variety of methods have been proposed in evaluating cybersecurity visualizations, ranging from quantitatively scoring network security visualizations based on data source categorization [90] to qualitative classification based on task analysis [89], but they often require complicated formulas or potentially biased categorization that is dependent on the evaluator. Furthermore, visualizations that score well on these types of frameworks have gone through multiple renditions and have evolved past their initial prototypes; however, one of the important aspects of creating effective visualization evaluations, that is often ignored, is a use case and subsequent task area that incorporates users throughout the entire evaluation lifecycle [25, 88].

With this in mind, I decided to use Endsley's "3 Level" model of situation awareness [4] to architect my study as Riverside was designed with the intent of providing high level cyber-SA to a user. D'Amico et al. walk through a general application of cyber-SA defense [36], which states that perception requires an analyst to take in the initial information and decide whether to escalate or ignore the initial data, comprehension occurs when an analyst combines their perception with previous mental models to correlate the potential escalation and trace suspicious traffic through a network, and last, projection allows an analyst to take the current incident and correlate it with common attack patterns or threats at a community level to take action enterprise-wide. Tyworth

et al. use the example of an analyst who first acknowledges an IDS alert, determines whether it's benign or malicious, and finally takes action on that alert based on their previous decisions [5].

5.2 Methods

Using the four user-centered information visualization evaluation guidelines given by Freitas et al. [25], I chose the context for my study to be an incident response scenario. My ideal user persona was a cybersecurity analyst [14], and I referred to the EEVi model [89] to guide the tasks that participants were asked to complete. I focused largely on SA and correlation analysis for incident response tasks to guide my scenario, as those are the two of the primary use cases for Riverside, as mentioned in Chapter 4. Last, I conducted this evaluation upon completion of a working prototype whose development was guided by user input from inception.

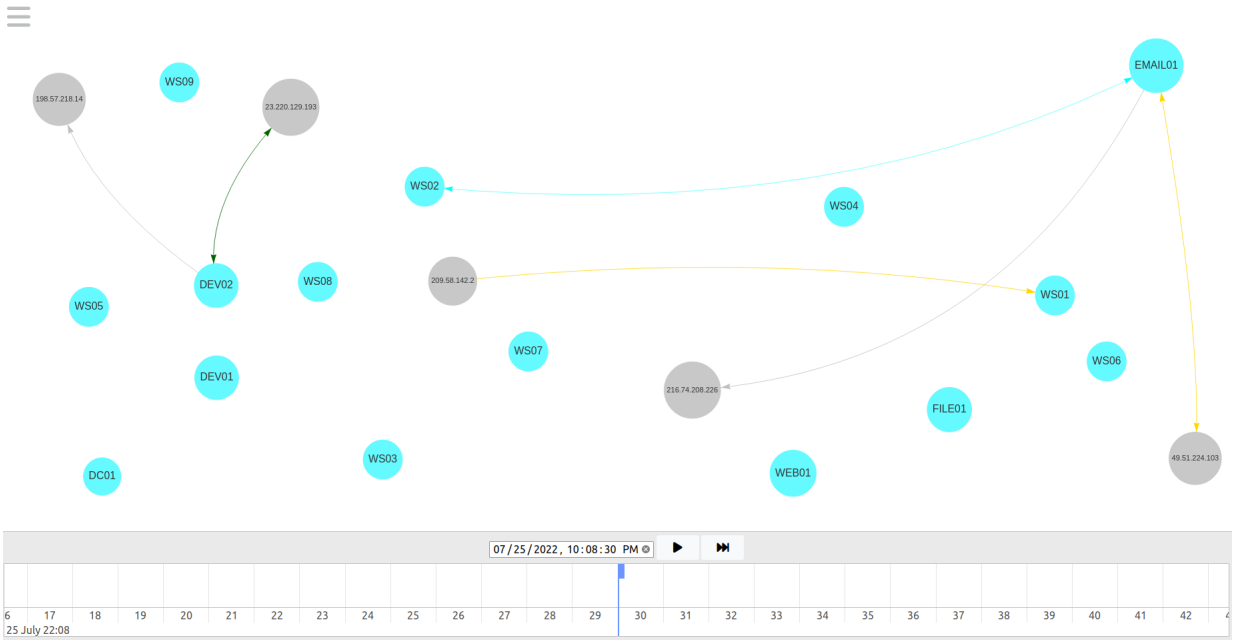

Figure 5.1: **Initial Riverside visualization for usability scenario.** This is the initial visualization users will see when they navigate to the time reported in the scenario prompt. WS02 is the reported workstation and at this time, it is currently communicating with EMAIL01.

5.2.1 Scenario

My scenario breaks down Endsley's 3-Level model to three levels of cyber-SA as follows:

1. **Perception (Level 1 SA):** A security analyst is told there is an alert on a host but only given a timestamp from the help desk who received the call from a user on the network. The analyst must determine if this activity warrants further investigation into the rest of the network.

2. **Comprehension (Level 2 SA):** After looking through all of data provided, analysts must determine if the attack chain is present on the network using Riverside.

3. **Projection (Level 3 SA):** Finally, the analyst must decide on a COA based on their findings.

I deployed Riverside on an internet accessible Digital Ocean droplet with pre-captured network data to perform my usability study. I used a a Qualtrics form to capture user responses and feedback. I created and deployed a small test network of 15 hosts using containerization and custom bash scripts to generate traffic [1]. Using these scripts, I simulated a realistic attack on the network, loosely modeled after the tactics, techniques, and procedures (TTPs) outlined in the MITRE ATT&CK framework [115] as seen in Table 5.1. This ensured accurate attack timing across all participant sessions to establish a consistent baseline for every participant trial.

[1] The source code used for building and orchestrating my attack scenario is located here: `https://github.com/artemis19/riverside_scenario`

Table 5.1: **Scenario components mapped to MITRE ATT&CK Techniques.** This shows the general category and technique as labeled in the MITRE ATT&CK framework mapped to the component of the scenario that uses similar TTPs.

MITRE Category	Scenario Component
Initial Access: External Remote Services	Brute-force RDP credentials for externally-facing email server
Command and Control: Application Layer Protocol (DNS)	Use of DNS for C2 beacons to communicate with C2 server
Lateral Movement: Exploitation of Remote Services	Use of *psexec* over SMB to move laterally throug thenetwork
Exfiltration: Exfiltration Over Alternative Protocol	Use of FTP to initiate data exfiltration from FILE01 out through EMAIL01 to a remote host

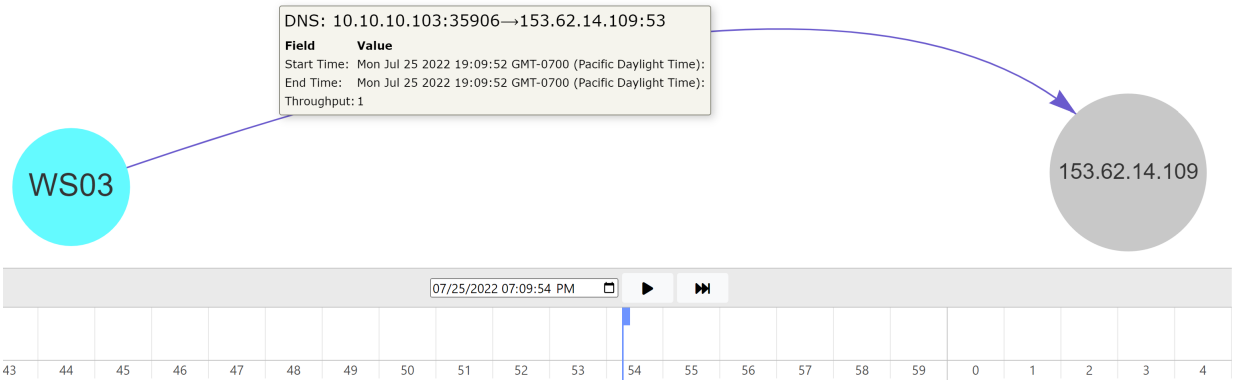

Figure 5.2: **C2 beaconing in scenario.** The remote host, 153.62.14.109, is the C2 server that establishes beacons on each compromised host over DNS.

Scenario Attack

The attack chain begins with RDP-bruteforcing a login to the email server, EMAIL01, as initial access. A second remote host acts as a command & control server (C2) and establishes a beacon on EMAIL01 over DNS. The "attacker" then performs lateral movement over SMB to WS02 then to WS03 and finally FILE01. In between the lateral movement, the C2 server establishes beacons on each compromised host, all over DNS as seen in Figure 5.2. The last piece of the attack chain was data exfiltration, which originated from FILE01, flowed through

Figure 5.3: **Data exfiltration in scenario.** The initial data exfiltration starts from FILE01 as is pushed through EMAIL01, the attacker's internal pivot point, and out to the malicious remote host, 153.89.194.75.

EMAIL01, and was exfiltrated from the network to a third remote host part of which can be seen in Figure 5.3. All of the remote hosts used the same first octet, "153," for their IP addresses.

Incident Response Use Case

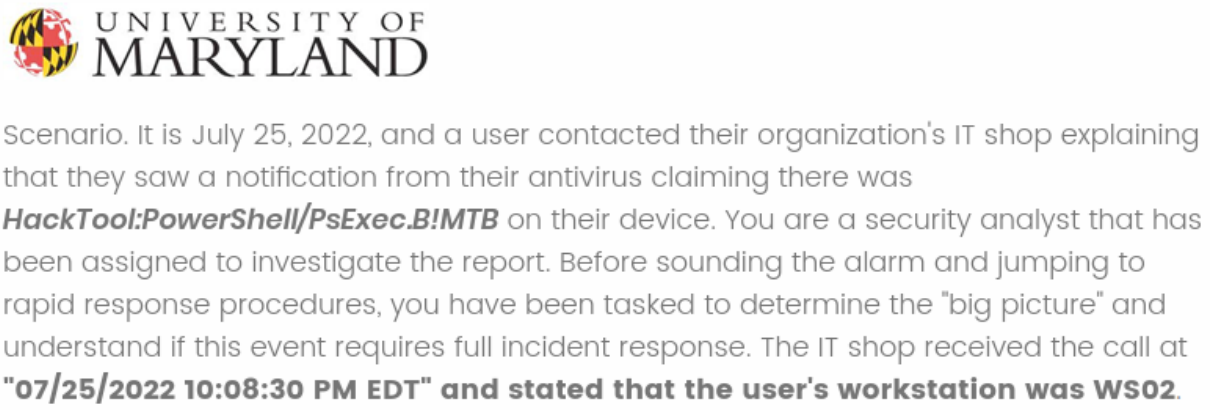

Figure 5.4: **Usability scenario prompt.** This is the prompt users are given when they begin the usability scenario. They are then asked to complete a variety of tasks and questions based off of this initial alert and their findings throughout.

The scenario starts with alerting the analyst that there was suspicious activity on a user's workstation which was reported to the IT shop as described in Figure 5.4. It is then their job to use the Riverside tool to investigate the network traffic around the time of the alert and determine the full attack chain by answering questions and performing tasks along the way. The initial visual participants see is shown in Figure 5.1. The scenario ends asking participants to summarize the entire attack chain and state what they would do next as as part of their processes and procedures as a security analyst. In addition to the required tasks, I had users interact with certain parts of Riverside's design, like node customization, to show how Riverside can help users track certain hosts during an incident investigation. The full set of scenario tasks and questions can be seen in Appendix B.1.

5.2.2 Participant Recruitment

I recruited participants using snowball sampling through university IT departments and personal networks. I wanted to ensure my participants represented real-world users of Riverside, so recruited participants were security professionals who perform perform network security or IR tasks in some capacity. I recruited participants during August 2022 and completed all 10 usability trials throughout August and September 2022. I refer to usability participants using "S" and their associated identification number, and questions using "Q" and the respective task or question number.

Table 5.2: **Usability Study Participant Demographics.** I show, for each participant, their gender, years of experience in the security field, age, and highest completed level of education (high school, bachelor's, master's, or doctorate).

S#	Gender	Years of Experience	Age	Education
1	Male	15	42	MS
Continued on next page				

Table 5.2 – continued from previous page

S#	Gender	Years of Experience	Age	Education
2	Female	12	40	BS
3	Female	15	43	MS
4	Male	4	21	HS
5	Male	5	26	BS
6	Female	1	29	BS
7	Male	17	50	PhD
8	Male	6	21	HS
9	Male	6	27	PhD
10	Male	9	33	BS

Table 5.2 contains demographics for my 10 participants, including gender, years of experience in the security industry, age, and highest level of education completed. My participants were fairly well educated with an average age of 33. Majority of my participants identified as male with an average professional experience of 9 years in the security field.

5.2.3 Usability Trials

I used a script so that each participant was verbally provided the same overview. The Qualtrics form contained a 6 minute instructional video showcasing Riverside's features along with a PDF overview that participants could download and reference for the remainder of the studystudy, as seen in Appendix B.2. I told participants I could not answer any questions about the scenario or the tasks they were required to perform. Participants could ask questions if they had any usability issues with the tool that prevented them from performing the required tasks.

5.3 Results

My analysis was largely qualitative in nature as this was a formative usability test to determine any usability issues and collect feedback [24] for the Riverside tool. I do showcase

quantitative results for task completion and semi-quantitative analysis of participant's achieved cyber-SA. Additionally, I present the quantitative results of my 5-point Likert scale focused on providing feedback from participants based on Riverside's ease of use, their experience, and Riverside's efficiency during their trial.

5.3.1 Quantitative Results

I started timing participants once they began the scenario portion of the study and stopped just before they reached the feedback questions. I had two prompts in the survey: one that alerted the user to let me know once they had started, and one that let them know when they completed their scenario questions and tasks. Table 5.3 contains participant times with the average time being 46:53 (min:sec). Participant trial times ranged from 30 to 60 minutes, with one participant who was time-capped at 60 minutes and did not finish the scenario(*). Any participant that is listed as having a 60 minute trial time did complete the scenario and took the entire 60 minutes.

Table 5.3: **Usability Study Trial Completion Times.** I show, for each participant the time it took them to complete all of the scenario questions or the last question that a participant completed if they were time-capped at 60 minutes(*).

S#	Time (min:sec)
1	40:13
2	37:48
3	41:50
4	33:43
5	50:06
6	60:00
7	3*
8	54:58
9	30:09
10	60:00

I used binary scores of 0 or 1 to determine whether a participant successfully completed the

required task or not. All of the quantitative tasks were used to calculate completion rates, which included Q1, Q4, Q6-7, and Q9-14.[2] The questions not included were qualitative in nature, such as asking a participant what color or shape they chose for a specific host during the scenario. Due to my small sample size (n=10), I chose to use an adjusted, or modified, Wald method with a 95% confidence level to analyze task completion rates [24]. Additionally, 9 participants are included when calculating completion rates for every question except Q1(*), as S7 did not reach any of the tasks or questions past Q3 as shown in Table 5.3. My task completion rates and confidence intervals are displayed numerically in Table 5.4 and visually in Figure 5.5.[3]

Table 5.4: **Usability Task Completion Rates.** This table contains the mean of the successful completion rate (p') for each quantitative task and the associated confidence intervals using an adjusted-Wald method with a 95% confidence level.

Question #	Mean (%)	Lower Interval (%)	Upper Interval (%)
1*	72	48	95
4	38	12	65
6	46	19	73
7	31	5	56
9	31	5	56
10	62	35	88
11	46	19	73
12	46	19	73
13	54	27	81
14	62	35	88

Overall, participants successfully completed 49% of the tasks. Q1, Q10, and Q14 had the highest completion rate, while Q4, Q7, and Q9 had the lowest with the remainder of tasks hovering just around or under 50% for all participants. The initial RDP compromise of EMAIL01 (Q4) was correctly identified 38% of the time, while participants accurately stated that EMAIL01

[2]For tasks that involved entering a timestamp, I allowed a $+/-3$ second buffer on my answer key, due to the 2 second data batching. This was done due to the differences of when the visualization actually displayed an entity versus what the "Created At" timestamp showed when hovering over a component.

[3]The mean reported in my results is the adjusted mean, or p', calculated using the adjusted-Wald method. The use of p here is completely distinct from p-values used to measure statistical significance.

was the first host compromised in the network (Q6) 46% of the time. The lateral movement of SMB across the network (Q10) was correctly recognized by 62% of participants. The C2 host (Q7) communicating over DNS (Q9) was only observed 31% of the time.

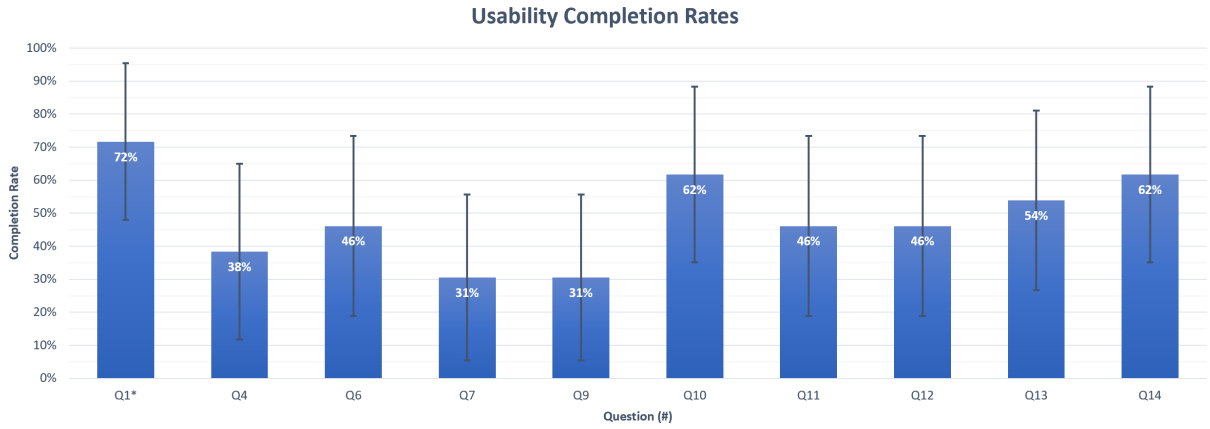

Figure 5.5: **Usability Task Completion Rates.** Each task's successful completion mean (p') is shown in the corresponding bars by question number (blue). The error bars (black) represent the lower and upper bounds of the confidence interval calculated with a 95% confidence level using an adjusted-Wald method.

The data exfiltration portion of the attack was spread out across 3 questions. Participants were 46% successful in identifying when and where the exfiltration occurred (Q11) and the other internal hosts involved (Q12), while 54% of participants correctly observed when the third suspicious remote host had data exfiltrated (Q13). The second-highest completion rate, tied with Q10, was Q14 where 62% of participants successfully identified the compromised hosts. While these numbers don't show significant results, they provide a benchmark for future work with Riverside as this was a formative study meant to determine the initial usability of Riverside.

5.3.2 Cybersecurity Situation Awareness Results

I used three questions to qualitatively assess the cyber-SA that Riverside provided to participants as described above. Q3, Q15, and Q16 were used to assess the three levels, respectively.

I performed semiquantitative analysis to showcase the relative cyber-SA that my participant pool achieved for certain levels, using key points from my answer key to determine whether a participant completely understood the attack chain or not. I generated the following answers to compare participant responses to:

1. Perception (Q3): Participants should **identify** the continuous SMB connections across the network around this time seems unusual or that the initial alert is indicative of something inherently suspicious between WS02 and EMAIL01 and should be investigated further.

2. Comprehension (Q15): Participants should have **analyzed** all relevant traffic and seen that the attacker likely bruteforced RDP credentials for the email server, starting around 10:07:17 PM, ending around 10:07:59 PM from the 153.17.52.201 IP address. The attacker then used 153.62.14.109 as their C2 server to maintain callbacks over DNS after moving laterally to a new host on the network over SMB. They accessed EMAIL01 initially, then moved to WS02, then WS03, and finally FILE01. They then exfiltrated data from FILE01 through EMAIL01 out to the 153.89.194.75 address, starting around 10:11:10 PM, ending around 10:12:20 PM.

3. Projection (Q16): Participants should **recommend** something along the lines of correlate IOCs across the rest of the network, block IPs across network, isolate infected hosts, or look up threat intelligence reports about IOCs to see what other TTPs could be present.

Perception

For Q3, participants had varying responses, but overall, 70% of participants (n=10) did mention something about the continued SMB communication between WS02 and EMAIL01 as a

reason to investigate further. Some participants, like S4, were incorrect in their initial assessment and mentioned what they thought was a portion of the attack chain since *"at 10:08:37 PM [EDT] WS02 is connected to by a new external host (185.199.111.133) which also connects to a number of other internal hosts all at the same time."* S6 even went on to mention *DNS requests outbound toward remote hosts which looks like non-normal behavior*, correctly identifying the C2 beacon for WS02, while S10 had identified the anomalous traffic with EMAIL01, specifically the *"RDP traffic with some random external IP and also making an SMB connection to a workstation."* Even though S7 did not complete the entire scenario, they did notice some suspicious activity which they mentioned in Q3, such as the *"large amount of traffic between WS02 and EMAIL01 shortly after antivirus alert," and the extended communication "between WS02 and WS03 after a bit longer after antivirus alert."*

Comprehension

89% of participants (n=9) identified WS02 and EMAIL01 as two of the four hosts that were compromised and thus mentioned those in their response for Q15, but only two participants successfully identified every component of the attack chain. Both accurately mentioned:

- Initial remote access to EMAIL01 over RDP
- All three of the correct suspicious remote hosts
- Lateral movement over SMB
- C2 beacons over DNS
- Data exfiltration that occurred from FILE01 through EMAIL01

- Four compromised hosts (EMAIL01, WS02, WS03, FILE01)
- Applicable timestamps for each portion of the attack chain

Four participants only missed one component of the attack chain. S8 identified everything in the attack chain except the initial RDP compromise and thought that *"some TCP-based spray occured with 185.199.111.133 at 7/25/22 10:08:08 PM [EDT],"* was the initial compromise since it occurred just before the SMB communication between WS02 and EMAIL01. Participants S5, S6, and S9 only missed identifying the correct C2 server and subsequent beacons over DNS. An example of one of the incorrectly perceived attack chains was given by S3 who believed the initial compromise was a phishing email to EMAIL01, but then claimed the *"lateral movement originated from 185.199.111.133 [remote host] across several hosts, starting with the DC01 server at the same time, 22:08:36 [EDT]."*

Projection

Most participants suggested appropriate recommendations for Q16, but dependent on their previous answers, they may have suggested that those actions be taken on hosts that weren't necessarily compromised, such as S1 who said *"WS02, DEV01, and DC01 all need to be checked and sanitized. Also an incident report needs to be drafted and a damage assessment done."* In this instance, hosts that were seemingly unaffected would have been taken offline, like DEV01 and DC01, but appropriate remediation such as isolation and a damage assessment were mentioned. S2 stressed the isolation of EMAIL01 as it was the first to be compromised, while S8 stated that *"there were several indicators of suspicious activity chained together that strongly suggest a breach, pivot and exfiltration occurred"* and recommended contacting an IR team for a more

in-depth investigation. Even though S3 did not correctly identify the compromised hosts, they did suggest useful host forensics techniques such as *"taking the $MFT"* and looking for *"file system changes that matched the timeline of the attack to see if there are other integrity issues"* that would have helped them confirm their findings.

5.3.3 Likert Scale Feedback

I gathered feedback using a 5-point Likert scale for Riverside's ease of use, participant's experience using Riverside, and Riverside's efficiency to complete the specified tasks. Figures 5.6, 5.7, and 5.8 contain the results from my feedback questions where participants could respond with "Strongly Disagree," "Somewhat Disagree," "Neither Agree nor Disagree," "Somewhat Agree," or "Strongly Agree" for each. Overall, participant responses were overwhelmingly positive, and all participants (n=10) provided responses for the feedback questions.

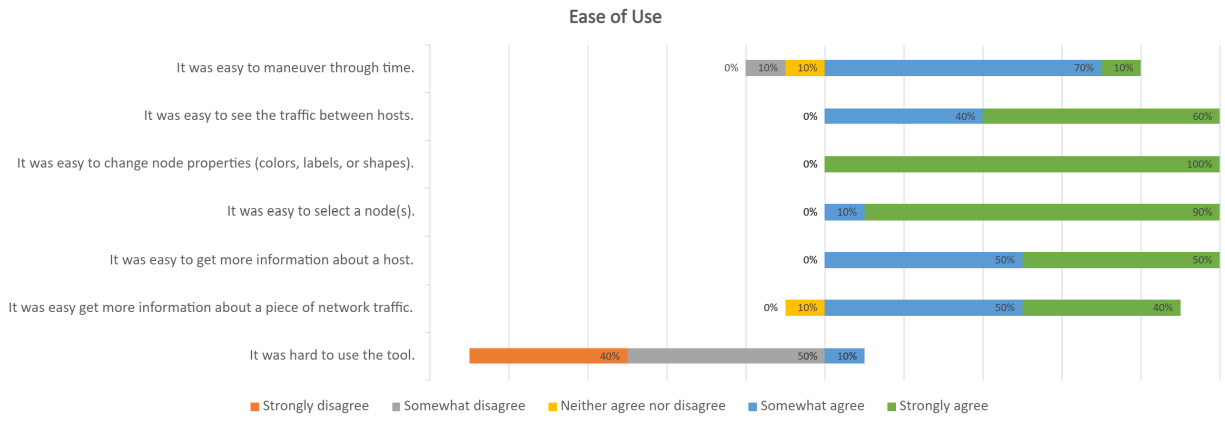

Figure 5.6: **Usability Study Ease of Use Feedback.** Feedback focused on participant's ability to use Riverside, broken down by the components that make up Riverside.

These questions focused on getting feedback about how easy Riverside was for participants

to use, and if the functionality of certain components was straightforward. Ease of use feedback was an important initial assessments as one of the goals for Riverside was to make a simple tool that security professionals could easily operate. An important result to note is that for all of the positively-framed questions, at least 80% of participants answered "Somewhat agree" or "Strongly agree," meaning that participants felt they could easily see traffic between hosts, change node properties, and easily see information about hosts and network communication. For the one negatively-framed question, 90% of participants disagreed that the tool was hard to use. The one category that had slightly varying responses was "It was easy to maneuver through time" where two participants answered "Neither agree nor disagree" and "Somewhat disagree," respectively.

Experience

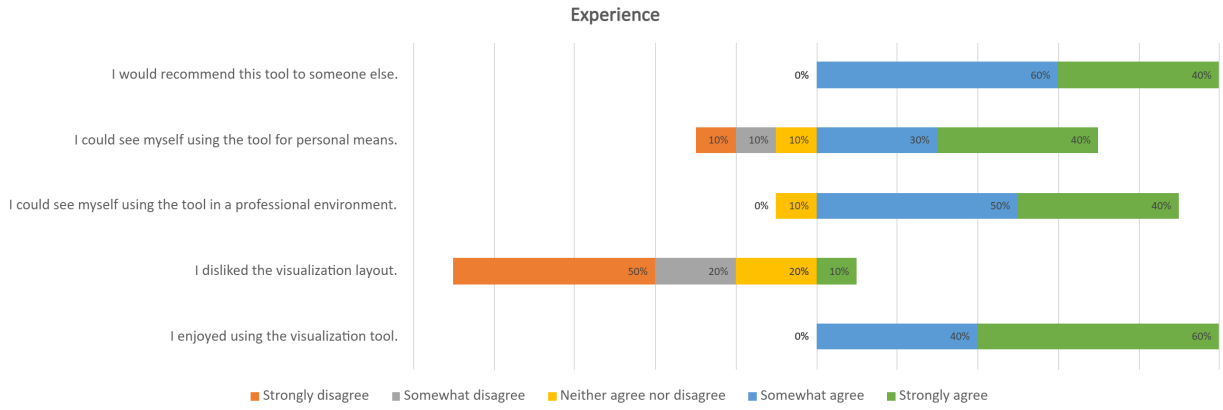

Figure 5.7: **Usability Study Experience Feedback.** Feedback results focusing on participant's experience using Riverside while working through the provided scenario.

I also wanted to know about participant's experience using Riverside in this specific application, and if they would use it again or even recommend it to a fellow colleague. One of the most important takeaways is that 100% of my participants said they would recommend this tool to someone else, with 70% saying they'd use Riverside personally and 90% that would use it in

a professional capacity. Only one participant strongly disliked the visualization layout with 70% agreeing that they enjoyed it, but 20% didn't have a preference one way or the other. Nevertheless, all of my participants said they enjoyed the functionality of Riverside when working through the scenario.

Efficiency

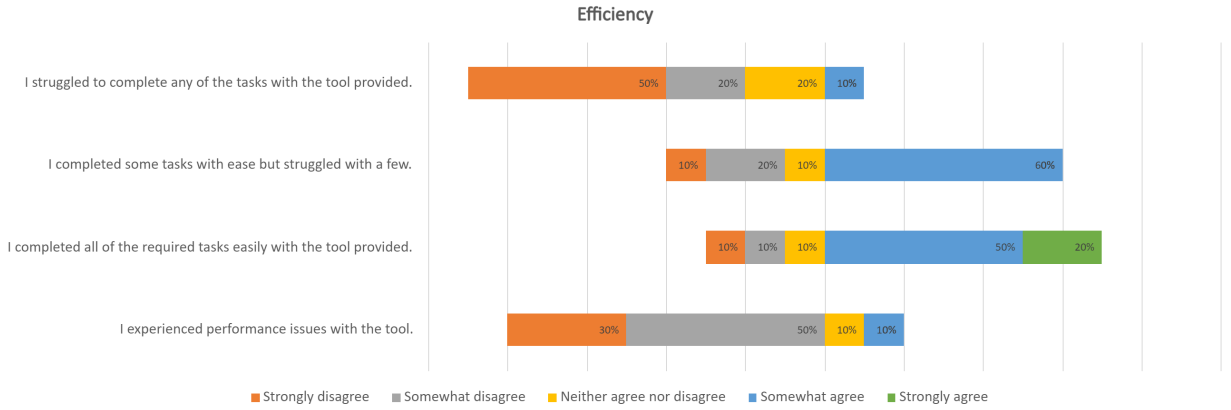

Figure 5.8: **Usability Study Efficiency Feedback.** Feedback results focused on getting feedback on participant's perceived efficiency using Riverside.

These feedback questions were focused on participant's perceived ability to complete the tasks in the scenario. Some key results to highlight are that 70% of participants said they did not struggle to complete any of the tasks provided and felt they were able to work through the scenario rather easily. Majority of participants felt that they struggled with a few tasks, which is to be expected when presenting users with a brand new tool. Additionally, 80% of participants did not experience performance issues with the tool in a way that meaningful impacted their ability to complete the tasks.

5.3.4 Usability Issues

There were some usability issues that either slowed participants down when getting started or caused some speed bumps along the way, but participants said there was nothing inherently wrong with Riverside that prevented them from working through the tasks at hand. Some participants did ask questions during the scenario to clarify the usability of a specific component or to make sure that the tool was behaving as intended. Some participants were initially confused on whether they were supposed to navigate anywhere other than the initial time provided, at which point I confirmed they could, allowing them to continue with their tasks. A few participants struggled to navigate to the initial alert time, in which case I had them confirm what time they had typed in, and in most cases, the participant had simply typed the date in wrong or forgot to specify "PM" and was looking at the correct time but in "AM."

Some of the basic usability issues that participants noted was a timezone specifier, because currently the timeline adjusts to local time, but I had the times in the scenario specified as Eastern Daylight Time, causing participants to have to adjust the times mentally depending on their time zone and navigate accordingly. Additionally, a couple participants noticed that if a specific piece of network communication was repeated later on, as in the same two hosts communicating over the same protocol and ports, the "Created At" timestamp would show the later time rather than the communication currently happening, which was correctly shown in the "Updated At" timestamp. In this instance, I told participants to go off of the current time specified on the timeline, and they were then able to continue. Last, when some participants were first starting to use the tool and experimenting with zooming and spanning the canvas, they lost where the original nodes were located and were forced to refresh the page.

5.3.5 Qualitative Feedback

After participants had completed all of the Likert scale feedback questions, I allowed them to provide both written and verbal feedback. Most participants expressed incredibly positive feedback, and a common thread across all participants was that they'd never used a visualization that presented the information the way Riverside did, particularly with the timeline component for navigation. P4 said *"it would be awesome if we could feed our insights into a visualization tool like [Riverside],"* while P5 appreciated that information was collected beyond just internal hosts, unlike other security-oriented visual analysis tools they've used.

P2 said the *"biggest issue [they] had was scrubbing through the time, at some points I completely lost the time window in question (I zoomed out way too far) and at other times scrolling along the visualization was "jumpy" and I couldn't get it to settle where I wanted it."* Participants that experienced this issue suggested that a "home" button for spatial reference or a timeline indicator would have assisted them in reorienting the canvas correctly. Many participants also noted that if they changed a node property early on in the scenario, it took them a bit to find it again. P9 and P7 both said they would have appreciated consistent node placement that provided insight into the logical setup of the network because in Riverside's current state it *"it was not clear how node and edge placement were handled"*. One of the most commonly requested features across participants was the ability to pause the timeline immediately upon dragging the blue cursor instead of having to pause and then drag, especially if they had seen something of interest and wanted to "scroll back" to look closer at it. Others desired a way to show the true source and destination, as it would have helped with their understanding of which host initiated certain communications. P7 thought that a legend for the edge colors would be helpful so that users didn't

have to *"mentally do the mapping"* of the associated ports and protocols.

Participants verbally echoed what the "Ease of Use" feedback questions showed in that they could see themselves using this tool again, and that Riverside was very *"easy to use"* and *"intuitive."* P1 had an interesting recommendation of being able to export a video from specified times, so they could include it in a presentation for superiors. Additionally, several participants specifically mentioned how that they could see Riverside being useful for correlation of alerts in tandem with other tools or incident investigation mechanisms to *"get the big picture view of what's going on."* Some of the "quality of life" feedback was having the ability to copy-and-paste times, node metadata, and traffic communication, so that they could easily perform lookups in other capabilities. Participants also said that having a rewind button in addition to a fast-forward option would have been helpful, as well as being able to set the playback speed of the timeline.

5.4 Discussions and Limitations

I will now discuss the results previously mentioned above by explaining why certain outcomes possibly occurred and generalizing certain results that can be applied to future work with cybersecurity visualizations.

5.4.1 Explaining the Results

The hardest tasks for participants appeared to be Q4, Q7, and Q9, all of which were below 40%. These questions correlate with identifying the initial malicious host, the C2 host, and the protocol used for beaconing. I think these questions gave participants issues because Q4 required participants to go back in time before the initial alert time, which might have caused

some confusion. Additionally, I don't know if all participants noticed that EMAIL01 had an external interface, which was meant to simulate an externally-facing email server, possible web mail. If they had noticed this, then maybe the initial RDP communication would have stood out compared to other remote hosts communicating with EMAIL01, as it was only the continued connection throughout the entire scenario between EMAIL01 and a remote host.

Since Q7 and Q9 were directly related, it should make sense that they had the same completion rate. Additionally, the C2 beacon over DNS was not as long and "obvious" as the SMB communication, so it was probably harder for participants to narrow in on that traffic, especially when that communication first happened. A common answer that participants provided as an alternative C2 remote host (Q7) was 185.199.111.133, because it communicated with a large portion of the network over ICMP shortly after the initial alert provided in the scenario. In hindsight, that is arguably suspicious traffic, but it also only happened once and was not something that occurred after the successive lateral movement over SMB. Q10 and Q14 were the second-highest completion rates. This shows that if participants were able to determine SMB as the "indicator" of lateral movement, then they were also able to correctly identify the compromised hosts on the network.

For the overall cyber-SA achieved, the average task completion was on par with the percentage of participants who identified the correct attack chain (Q15) and therefore recommended actions to be taken on the correct hosts within the network (Q16). Almost all participants acknowledged that the initial alert and corresponding SMB traffic led itself to further investigation, but some participants did get themselves off track to start by focusing on something that was potentially not part of the attack chain at all. Participants that did this seemed to have tunnel vision, which can be common among security analysts, especially when they have no other ca-

pabilities or personnel to confirm their findings. Something I observed was that some of the participant's answers for the individual questions didn't always match what they determined as the attack chain at the end of the exercise. This can be seen with P5 who incorrectly answered Q4 as *"185.199.111.133 (10:08:08 PM [EDT])"*, but then correctly identified the three suspicious remote hosts and the times that MCA likely occurred on the network when answering Q15 and Q16. This could have been because they didn't want to go and correct their answers or because they were running out of time and didn't have the time to go back.

For participant's overall success using Riverside with the provided scenario, I think it was hard to pinpoint anything specific that caused participant's performance to be sub-par. P1 expressed the sentiment that *"I'm sure I got stuff wrong because when there's an attack, everything seems suspicious!"* Other participants said that not having that baseline knowledge of the environment they were working in, such as a network map or asset list, was a hindrance. P4 stated that *"it's hard to know what I should be looking for because this isn't my environment, and I don't have logs to correlate."* I think my results showcase that it's hard to develop all-encompassing visualization tools, and security is a dynamic field that lends itself to having a breadth of capabilities. Last, even though the feedback largely positive, participants completed all of the Likert scale feedback questions without knowing how well they'd performed, so while most participants answered positively with these questions, some might have changed their responses had they known their performance results. In particular, the results for participant's efficiency were more mixed than the other feedback categories since participants answered these questions without knowing their "success" for the scenario.

5.4.2 Generalizing the Results

I will now discuss how my results generalize to the broader fields of information visualization and cybersecurity. While my sample size (n=10) was small, my participant pool was representative of the cybersecurity community. My participants were an average of 33 years old, 70% male, and 30% female, and according to a 2022 survey, the average age of a cybersecurity professional is 42 years old with 78% identifying as male and 22% identifying as female [116]. This supports that my results should generalize to larger populations, as I wouldn't need to use an adjusted method with a larger population, and task confidence intervals would actually become narrower and more precise. Additionally, I developed a realistic incident response scenario and tied it to a use case that security analysts have to perform on a regular basis: identification and analysis of anomalous network activity. I also constructed my attack chain based on industry standards of common attacker TTPs. While this was a controlled environment, my results showcase the potential for using visuals to aid security analysts in performing correlation analysis, and the need for using realistic scenarios when constructing usability studies.

Additionally, I posit that the compellingly positive feedback on Riverside's design and visualization piece supports my bottom-up approach and can be applied to future cybersecurity visualization development. Maintaining a "simple" and "easy to use" interface is easier for users to digest, compared to the complicated and crowded visuals commonly used in cybersecurity. Security analysts need immediate visuals when balancing numerous alerts and capabilities, which reinforces the need to "display all facets of an environment without user intervention." Moreover, allowing users to characterize the visualization makes them feel as though it's theirs and further motivates tool adoption. Going back to some of the capability gaps identified by prior research,

participants were able to correctly identify the malicious lateral movement over SMB, 62% of the time, and 70% of participants discussed suspicious SMB connections across the environment when analyzing their perception of the initial network activity. Additionally, 60% of my participants were able to correctly identify all or almost all of the attack chain with just an initial prototype of Riverside, showcasing its potential for augmenting cyber-SA, particularly during a critical event such as an incident response case.

All of the above points underscore the importance of balancing automation with user input, especially when dealing with dynamic data in a technical field. Riverside provides a dynamic visual of a network's behavior to provide analysts that missing internal baseline. Additionally it ties in the external component of an environment that is necessary to ensuring cyber-SA across the entire domain of an analyst's responsibility. Additionally, I show the possibilities for using information visualizations to augment cybersecurity capabilities, specifically with the use cases of situational awareness and correlation analysis. Furthermore, incorporating users into the entire lifecycle of a tool will lead to acceptance within the community and increase the chances of integration into a user's processes and procedures.

5.4.3 Limitations

My participant size was very small, and since I used snowball sampling, I acknowledge that this could have limited my results in establishing a realistic benchmark; however, I did manage to recruit a broad range of participants in terms of background and experience. Additionally, participants were probably keyed in that something anomalous was occurring on the network whereas realistically, security analysts may just be keeping an eye on things without necessarily

hunting for MCA.

Additionally, my scenario was developed entirely by me with input from an expert in visualization but not in security, so question wording or ordering could have impacted my results. For example, originally the answer to Q11 was FILE01 in my answer key, but after reading through some of the participant responses, I noticed that many people answered EMAIL01. When I reread the question, I realized it was slightly ambiguous through the use of the word "final," which was originally intended to be the final host compromised, but technically EMAIL01 is the final host accessed by the attacker before data is exfiltrated from the network. I then decided to accept either host with a correct timestamp as correct answers.

I also use pre-captured network data and did not have participants use the real-time functionality of Riverside. This was done to maintain a repeatable scenario across all participants for accurate analysis. Additionally, an analyst would normally have other tools and personnel at their disposal to confirm their findings, but I required them to complete their tasks alone and using only the Riverside visualization. While most of my analysis was rudimentary, it provides a first cut in quantitatively analyzing achieved cyber-SA using a cybersecurity visualization.

Chapter 6: Conclusion

6.1 Summary

In this thesis, I presented my analysis from interviews with 24 network and security professionals which resulted in 14 themes and 24 visualization features mapped to user actions. Using my interview results, I designed an initial prototype a network security visualization tool, Riverside. I conducted a formative usability study with Riverside to determine its usability, gather feedback on its current design, and showcase how Riverside can be used to assist security analysts in maintaining cyber-SA of their networks. I directly showcased the application of Riverside for two of the aforementioned use cases, big picture SA and incident correlation, through my usability study and direct feedback from participants. However, I make the argument that Riverside's architecture alone would assist with the third use case of asset management. Even in the event that an internal host is unaccounted for, as long as it communicates with something on the network that is monitored, Riverside will visualize the internal host as it communicates with a "known agent." This would provide an analyst the necessary information to deploy an agent to the unmonitored host and start building out proper asset inventory.

6.2 Future Work

In its current state, Riverside could become cluttered for larger networks, so techniques such as clustering could be used to assist with scalability [117] and spatial reference. Some of the more complicated features, such as frontend filtering or additional views, would need to be added for Riverside to be deployed in a true operational capacity, but more testing is needed before integrating those into the Riverside system. I would also like to add the ability for agents to have a criticality tag that would allow users to filter on types of hosts based on how "important" they are. Last, I want to incorporate node and timeline "staining," which would perform two functions: 1) add a tag to the timeline when a user interacts with a node and 2) stain a host that a user thinks is "suspicious" to track all other hosts that it communicates with by adding a stain. This would assist users in tracking their actions, as well as showcase the potential impact of a given network incident. Furthermore, I'd like to explore how event or hybrid staging could be incorporated into Riverside, as time-based staging has been shown to be less effective for online dynamic networks [118], even though it is the natural way to handle temporal data.

In the future, I'd like to conduct a survey presenting different mock-ups to get broader feedback on the overall design of Riverside. After incorporating feedback from this into Riverside, I'd like to perform an additional usability study that provides an analyst a traditional view of network traffic and alerts and compare it to quantitatively measure achieved cyber-SA, similar to the study performed by Giacobe [119]. I would then like to redo my usability study and compare my original baseline to the improved interface to see if any of the features mentioned above have a direct impact on the effectiveness of Riverside's visualization. Last, I want to deploy Riverside in an operational environment to complete a full usability assessment using the system usability

scale [120] to provide a quantitative measurement of usability.

6.3 Research Implications

At a high level, due to the technical nature of security, involving users in the initial designs and prototypes of a tool is paramount. This can be done in a variety of ways as shown through other research [14, 20, 92], but if you have direct access to security professionals, it is best to conduct user studies to directly align user needs and requirements to technical features and tasks. This not only invests users in the development lifecycle but ensures that the features incorporated into the visualization are directly related to real tasks that users will perform. Additionally, I posture that gathering input from adjacent fields will help identify components used by other professions that can be applied to cybersecurity visualization, similar to how I interviewed network administrators and NOC analysts. The themes and features presented in Chapter 3 showcase critical components that visualization developers should consider when designing cybersecurity visualizations.

Riverside's architecture was intentionally designed to limit required user intervention while supporting user-focused features. While Riverside's design itself received overwhelmingly positive feedback from the usability study, my process of involving real users in the design and evaluation process resulted in favorable responses and willingness from users to incorporate the tool into their daily processes. Additionally, participants felt that the user interface was simple, while still allowing them to interact with the visualization in ways that were meaningfully helpful for completing their tasks. Furthermore, despite the fact that Riverside was a prototype when evaluated, participants still managed to complete on average almost 50% of the usability tasks,

and 60% of participants identified all, if not almost all, of the correct attack chain. This highlights the importance of designing and evaluating with a specific use case in mind, so that a tool can be measured to its fullest potential.

My research showcases the gap that currently exists between visualization tool developers and security professionals, or the users, of cybersecurity visualizations. I provide several themes that should be used to guide future development of security visualizations while highlighting the necessity of incorporating user-centered techniques in the information visualization discipline, specifically those produced for highly technical fields. Security personnel should be considered stakeholders, and their input should be used to shape and mold the features and functionality of the visualization. One of the advantages in cybersecurity is defense in breadth and the ability to have overlapping capabilities that provide a variety of insights. Visualizations represent one type of capability. As highlighted throughout this research, security professionals are unlikely to rely on a singular capability to maintain SA of their environments. Visualizations are uniquely positioned to augment an area of security that lacks effective solutions.

Riverside's design fills the gap of low network visibility expressed by cybersecurity defenders, through a balanced visualization that provides analysts a big picture view of their network and data-driven insights to maintain cyber-SA across all aspects of their environment. I assert that Riverside's design can and should be used to augment current cybersecurity capabilities that already collect the necessary data. Throughout this thesis, I've shown the potential for using cybersecurity visualizations to supplement current capabilities by providing high level cyber-SA to analysts through the use cases of real-time situation awareness and correlation analysis. I have demonstrated that for cybersecurity visualizations to effectively provide cyber-SA, a bottom-up development approach, that is architected around end user actions and automatically visualizes

all aspects of an environment, is necessary.

Appendix A: Interview Study

I conducted semi-structured interviews with 24 participants. This section contains the recruitment messages, demographics questions, the full set of interview questions, and my codebook.

A.1 Interview Recruitment Messages

Below are some of the recruitment messages I used for the interview portion of my research.

(a) **Twitter Recruitment Post**

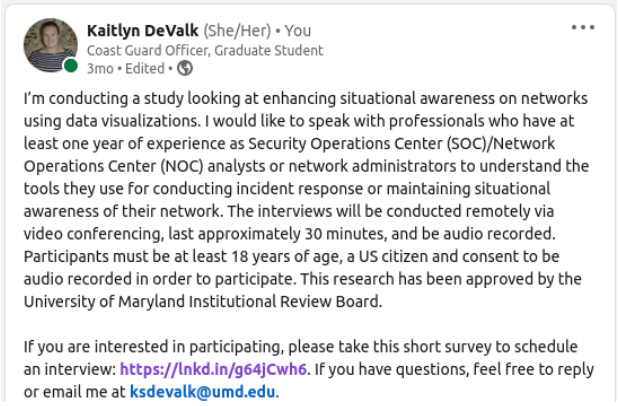

(b) **LinkedIn Recruitment Post**

Figure A.1: **Interview Recruitment Messages.** I posted on social media sites using the above messaging to recruit interview participants.

A.2 Questions

A.2.1 Demographics

These questions were included as part of the pre-screening survey.

1. How many years have you held a role as a network administrator and/or SOC/NOC analyst?
2. What gender do you most identify with?
3. What is your ethnicity?
4. What is your highest level of education?

A.2.2 Introduction

Hi, my name is Kaitlyn DeValk. Today, I will be conducting an interview with you to learn more about your experiences with network and security monitoring tools and understand what is necessary for you to be successful in your role. Have you read and signed the consent form provided to you in the pre-screening survey? (*If the consent form has already been received, skip the previous statement.*) You have the right to skip any questions that you are uncomfortable answering or refrain from answering anything that pertains to an aspect of your job that you cannot talk about. Please keep in mind that you have the right to withdraw from the study at any time, at which point I will discard any and all information I have collected from you.

A.2.3 Interview Questions

1. What area of industry do you currently work in?

2. What is your job title, and do you have a supervisory role?

3. Approximately how large is your current organization's network (endpoint count)?

4. How many analysts does your organization have (i.e. how large is your team) for maintaining situational awareness of your network?

5. On a scale of 1-5 (5 being excellent, 1 being poor), how would you rate the overall visibility that you have of your company or organization's network?

6. Do you have an accurate and up-to-date list of all devices/machines on your network?

    (a) (If anything is low (1-2) or "no" response) What do you think is needed to remedy these issues?

7. What types of tools have you used to perform network monitoring i.e. network captures, SIEM, EDR, DLP (Ex: Wireshark, Elastic, Cisco Prime, Solar Winds)?

8. Have you ever used a network visualization tool to assist in your job tasks? (Provide some examples of tools, if necessary)

    (a) If so, can you name them and give a short synopsis on what you liked or disliked?

9. Are you satisfied with the current tools your organization uses for maintaining situational awareness of your network?

    (a) Why or why not?

10. In your opinion, what is the most critical information or data that an analyst or administrator needs to obtain during an incident response investigation?

11. What user interface or layout would you envision for a network visualization tool and are there any key features that should be included? (Allow participants to draw a basic sketch or overview of their design or verbally share some features that they feel would be helpful.)

A.3 Tools Discussed

This is a list of all of the tools mentioned by participants when discussing network and security monitoring capabilities.

Table A.1: **Network Security Tools Mentioned by Participants**

Tool	Number of Times Mentioned
Acunetix	1
ArcSight	1
Arkime	1
Auvik	1
BlueCoat	1
Cacti	1
CarbonBlack	1
Chainsaw	1
Checkpoint	1
Cisco Firepower	1
Cisco Firesight	1
ConnectWise Automate	1
CrowdStrike	4
Dash	1
DataDog	1
Elastic Stack	6
Exabeam	1
F5 BipIP	1
FireEye NX	2
Gigamon	1
Google Stenographer	1
Grafana	1
Graylog	1
HP OpenView	2
LogRhythm	1
Microsoft Sentinel	2
Naggios	1

Continued on next page

Table A.1 – continued from previous page

Tool	Number of Times Mentioned
Nessus	1
NetCool	1
NetIQ	1
OpenNMS	1
OpManager	1
Palo Alto	1
QRadar	4
Qualys	1
Securonix	1
SentinelOne	2
Shodan	1
Snort	2
SolarWinds	5
Splunk	8
StellarCyber	1
Suricata	1
Tanium	1
Terracotta	1
TippingPoint	2
VMWare NSX	1
Windows Defender	1
Wireshark	1
Zabbix	1
Zeke	1

This is a list of all of the visualization tools mentioned by participants when discussing network and security visualization capabilities.

Table A.2: **Network Visualization Tools Mentioned by Participants**

Tool	Number of Times Mentioned
Arkime	2
d Auvik	1
BloodHound	2
Cacti	1
Cytoscape	1
DrawIO	1
Gephi	1
Grafana	1
Graphviz	1

Continued on next page

Table A.2 – continued from previous page

Tool	Number of Times Mentioned
Kibana	2
Microsoft Visio	5
MISP	1
NetWitness	1
QRadar dashboards	1
SolarWinds	3
Splunk dashboards	4
vRealize Network Insight	1
Wireshark	1
Zenmap	1

A.4 Codebook

I provide definitions for all of the codes used in my final codebook. These are the same codes located in my theme maps, as shown in Figures 3.3, 3.4, 3.5, and 3.6.

Table A.3: **Interview Codebook.** This showcases each code along with its definition and the number of times it was coded in my interview corpus.

Code	Definition	Count
Ability to communicate information effectively during an incident is paramount	Participants claimed that information flow among their teams was one of the most important parts of IR, whether that be the information being communicated or that people knew who to communicate it to	6
Ability to link incident to the network logs or data is an important capability	Ability to link alert or event on the network to the actual data internally is something participants stressed importance for in incident response/management	12
Access to top-of-the-line tools	Participants were satisfied with their capabilities because they have access to the best tools in the industry	2
Asset management is complex	Participants gave this as a reason for why they might not have an accurate list of devices or network map and that it's a complicated issue	4
Auditing would help with asset management	Participants expressed lack of processes or policy and communication as well as people following not those policies hurts asset management in their companies	5

Continued on next page

Table A.3 – continued from previous page

Code	Definition	Count
Average network situational awareness	Scores of 3 on visibility of network or where participant provided an answer of ”yes” with a caveat to having an updated list/map of devices on their network	9
Building a timeline can help with incident response	Participants said that being able to create and document a timeline of events from network data can help with incident management	2
Capabilities aren't being used to their capacity	Participants are somewhat satisfied with the insights they have but feel they could do more with their tools and either don't have the support, resources, or knowledge to do that	4
Capabilities provided do not meet needs of user	Participants worked organization where certain capabilities or tools were dictated by someone other than their team and that may not be what their specific team necessarily needs	3
Custom tools developed to meet needs	Participants said their company has developed custom in-house tools or capabilities to meet their needs or that they will do that when a tool they have isn't doing what they want	3
Data aggregation is important	Participants feel like at the end of the day, having the raw logs or data and having them in one place helps them when they need to investigate something	5
Dislikes inability to use tool outside of professional environment	Participants feel that tool is either too expensive or inaccessible for people to use outside of a corporation and therefore leads to inability to become proficient at it outside of work environment	3
Dislikes interactive portions of tools	Participants dislike having to activate tools to then click on pieces of a visualization, would rather just be able to interact with the nodes or visualization parts themselves and have the visualization update automatically from that	2
Dislikes requirements of additional resources or requirements for capability	Participants either disliked that a tool had additional software or hardware requirements to correctly function, making it cost more in money or time	3
Dynamic network mapping with host discovery can improve network visibility	Participants stated that they felt a dynamic network map or visualization would be useful for their job, particularly for identifying hosts on their network that may be untracked or no longer there	6

Continued on next page

Table A.3 – continued from previous page

Code	Definition	Count
Familiarity with environment is more important than any tool output	Participants felt that no tool is going to have the ”answer” and that analysts need to be familiar with their network environments which is paramount during an investigation	5
Graphical view provides insight into network topology	Participants that either have used graphical views of networks and appreciate being able to see the individual hosts or network from a big picture view or that didn't have that in their environments and wish they did	15
Great network situational awareness	4 or 5 in visibility or ”Yes” for up-to-date device list	17
Implementation of tools by others was done poorly	Either from an engineering perspective or from other analysts where participant couldn't then change the dashboard or tool layout or missed alerts and lower network visibility	2
Invest in dedicated capabilities for asset management	Participants described this either as tools that are meant to perform asset management or a person who owns the asset management component	1
Known unknowns are a blind spot in networks	Participants claimed that it was understood that there were probably hosts not being accounted for that were inadvertently discovered during an incident or that couldn't run endpoints (ex: IOT devices); usually due to network size, poor asset management, or environment requirements	6
Large network	Participant works on a large network (thousands of endpoints)	14
Long-term packet capture storage	Mechanism participant's corporation uses to maintain situational awareness on their network	1
Moderately-sized network	Size of participant's network is in the low thousands of endpoints	5
Monitoring and tool redundancy	Participants claimed that they either have tools or capabilities that overlap enough in areas that they feel they have no blind spots or that they have a backup capability	5
Multi-dimensional tool that provides multiple use cases	Participants use a particular tool for multiple functions within their network or security monitoring	2
Multiple teams or large for maintaining situational awareness	Participants work for corporation where there are multiple teams/divisions to manage their networks, usually include 24/7 monitoring	12

Continued on next page

Table A.3 – continued from previous page

Code	Definition	Count
Network visualizations are only one piece	For some participants, they care more about the forensics or individual host level, so the network visualization is something used but not their focus or main tool	3
Network visualizations should provide basic information that can be used to pivot	Participants felt that a visualization should provide basic data about a host or traffic to then easily pivot to another part of the visualization or tool when investigating something on the network	3
Network visualization tools can help explain things to others	Participants said that they like using visualizations to help explain things to upper-level management or to non-technical people	3
Poor network situational awareness	Scores of 1-2 or ”no” to visibility of network and having an updated list/map of devices on their network	8
Prefers 3D or rotational network diagrams	Participants said they prefer a 3D model or movable network diagram where they can see ”behind” something to understand the topology better	1
Prefers big picture layout to start with ability to dip deeper	Participants like the visualization piece to be the main focus to start and can have filters or pop-ups for more detail since they dislike how much data some of the tools have up front	8
Prefers combination of charts and graphics for visualizing data	Participants like a combination of charts and visualizations to showcase data	4
Prefers heat map style visualization to show network areas to start	Participants liked the idea of showing general network areas and then as you zoom or scale, it can show hosts in a certain part of the network	1
Prefers open-source tools	Participant likes open-source tooling since they can make modifications to it on their own, have more control over the functionality, and avoid large expenses (cost)	3
Prioritization of assets can help with asset management	Participants said that prioritization of the most important assets can help drastically improve asset management, especially during critical times like a network incident	1
Situational awareness on a network is never 100%	Participants claims that analysts should never be fully satisfied with the tools or network state	6
Small network	Size of participant’s network is in the hundreds of endpoints	5
Small team for monitoring	Participants work for corporation where it’s a singular team or group of people who manage the network	11

Continued on next page

Table A.3 – continued from previous page

Code	Definition	Count
Tool abstracts too much from user	Participants felt that tool didn't give user enough detail into the data, whether that was that inaccurate data was being fed in or that tool just didn't provide enough insight with no way to change that on their end	4
Tool aggregates network data automatically	Participants liked tools that put all of their data in one place so they could easily use it to make visualizations or graphs	3
Tool easily links network devices together	Participants liked that a tool they used showed how the devices on their network were connected and/or that they could provide some information about the connection between hosts	4
Tool is hard to use	Participants felt the tool was hard to learn how to use or complicated from personal experience or from watching others use it	7
Tool needs to be flexible to the environment it's deployed on	Participants claimed that either size or complexity of network make it hard to visualize with tools they've used, particularly because of too little or too much data that flows on their network	7
Tool presents too much information at the start	Participants said a dislike they had for tools was those that present a lot of information, especially on the initial screen which can lead to burn-out or complacency	3
Tool requires manual intervention from user	Participants described tools that required manual intervention from them to get alerts to trigger (workflows they have to build) or that the visuals need to be manually built out	9
Tool wasn't cross-platform compatible	Participants felt certain tools catered towards certain environments and didn't have good cross platform capabilities	1
Tools are only as good as the data fed into them	Participants made the statement that any tool is only as good as the data being fed into them which can affect the performance or story that the tool tells	3
Tools are used for something other than intended purpose	Participants described having to use tools for something they weren't intended to do because they lacked a certain capability, both positive and negative	4
Tools were tuned well for user's needs	Participants felt that they were currently satisfied because their tools or capabilities were tuned well to their goals or use case	3

Continued on next page

Table A.3 – continued from previous page

Code	Definition	Count
Understanding the attack vector is important for incident management	Participants said that understanding what an attack was targetting and if there is a threat actor involved what their TTPs are to prevent future attacks and to inform recovery decisions	3
Use scoping to deal with incidents	Participants said that it was imperative to understand the initial scope of an incident (i.e. where did this start, where is this on the network, what was compromised initially, etc) and then the extent/impact of the incident before taking further steps	11
Verbose logging can help with network visibility	Participants said they use very detailed logging to increase their visibility or situational awareness on their network	1
Visualization Features	This represents features participants mentioned when discussing specifics about designing a network visualization tool and will not be used for thematic analysis directly.	0
Ability to add criticality tag to hosts	This can be either pre-set or during an incident so analysts know which hosts are important to keep track of; can help keep track of assets or network segments that are more critical or important	5
Ability to classify or label network entities based on historical data	Either user or engine driven; remember classifications or labels analysts give traffic or connections so they don't have to keep tagging it repeatedly as good or bad and be able to lets analysts know when common "bad" things are happening	5
API-compatible to connect to other APIs or resources to correlate data	Being able to automate either searching for something like an IP address by connecting to other tool APIs or to automate pieces of the tool through its own API to save them time	2
Ability to correlate incident or case management alerts to traffic	Tag traffic or activity as an "incident" and have the incident management built into the tool and connect this to alerts	4
Ability to drill-down for more detailed data	"Less is more" up front with the ability to get more details as you click through or hover over components	11
Ability to filter data and save these filters if needed	Search for specific hostnames, protocols, IP addresses, internal/external connections OR automated filtering beforehand for traffic that is just "noise" in user's environment	8

Continued on next page

Table A.3 – continued from previous page

Code	Definition	Count
Ability to have different symbols for different types of hosts	Either through images, shapes, or patterns to differentiate entities other than through colors	3
Ability to have different views or dashboards for the visualization	Show data in different formats, whether that's tables or dashboards or different visualizations so people can have a choice on what they prefer to look at and have these views be customizable (ex: static vs dynamic flows)	10
Ability to have multiple environments or profiles	Able to load multiple network maps or environments, especially for multi-tenant environments, and have these be customizable by the user	2
Ability to highlight parts of the graph & show stats	Get stats on only some hosts or a given network segment	2
Ability to label nodes or hosts on graph	Either with things from the hosts themselves or for analysts to add notes	4
Ability to scale or move visualization	Screen is "infinite" so you can put what you want to see on the page or zoom in on certain parts	3
Ability to see metadata for hosts or machines that are unaccounted for	Being able to query/see information about a host that is possibly unaccounted for on the network in terms of asset management or endpoint detection but communicating with a host that is	1
Ability to see shared resources for an application or host	See links between hosts running a certain application or sharing network resources to understand the impact of an incident or traffic	3
Ability to share layout or display with other users	Saving or sharing visualization layout in standard format to import into another tool or to have another analyst pull it up on their screen	3
Ability to show amount of network communication visually	Change the size of the line connecting a host to another dependent on how many packets there are for a specific connection/communication	3
Ability to show general counts of interest to a user	Have some general numbers on the visualization initially (i.e. number of agents) to provide a high level "lay of the land" to start	1
Ability to show network from a geographic perspective	For global companies, having the ability to see where a host is in the world and be able to zoom in on a country, state, city, etc	2
Ability to show network segmentation	Show segments of the network so it acts like a dynamic "paper" network map to see how everything is connected	7
Ability to show snapshots in time of the network state	Show how traffic or the network has changed over time	2

Continued on next page

Table A.3 – continued from previous page

Code	Definition	Count
Ability to showcase common anomalous traffic to guide users to investigate	Have some sort of chart or pop-up of common anomalous traffic so that users have something to guide them while monitoring their network	1
Ability to use colors to convey some event to user	Use colors to convey network events or incidents (such as ”good” or ”bad”) so that when it comes up, analysts immediately know what it is	3
Ability to visualize network communication between hosts	Show data, protocols, or general communication between hosts, either internally or internal to external, visually (either statically or dynamically)	9
Ability to collapse or pop-out menus versus static option	If there are options for selecting things, make menus collapsible or only pop-ups so they don’t take real estate away from the visualization components	1

Appendix B: Usability Study

This contains the demographics questions, scenario tasks and questions, and the Riverside instructions.

B.1 Questions

The following questions and participant responses were captured in a Qualtrics form.

B.1.1 Demographics

1. How many years of experience do you have in the security field?
2. What gender do you identify with?
3. What is your age?
4. What is your highest completed level of education?

B.1.2 Scenario Tasks and Questions

Scenario: It is July 25, 2022, and a user contacted their organization's IT shop explaining that they saw a notification from their antivirus claiming there was *HackTool:PowerShell/PsExec.B!MTB* on their device. You are a security analyst that has been assigned to investigate the report. Before

sounding the alarm and jumping to rapid response procedures, you have been tasked to determine the big picture and understand if this event requires full incident response. The IT shop received the call at “07/25/2022 10:08:30 PM EDT” and stated that the user’s workstation was WS02.

1. Navigate to the time of the user’s reported event on the timeline. How many ”remote” or gray nodes are present on the visualization at this time?

2. Tag the internal host you are investigating by giving it a distinct color, shape, or label. Which option did you choose?

3. Look through the traffic around the time of the call. Why do you think this alert warrants further investigation on the network?

4. Using the network data provided to you, determine the first suspicious remote host IP address that communicates with an internal host and at what time. Enter the IP address and time below (HH:MM:SS AM/PM).

5. Mark the first suspicious host by giving it a distinct color, shape, or label. Which option did you choose?

6. What is the first host that was compromised in the network? Choose the host below.

7. Using the network data provided to you, determine the second suspicious remote host IP address that communicates with an internal host and at what time. Enter the IP address and time below (HH:MM:SS AM/PM).

8. Mark the second suspicious host by giving it a distinct color, shape, or label. Which option did you choose?

9. What protocol is the second suspicious host using to communicate with the internal network?

10. What protocol is the attacker using to move laterally through the internal network?

11. What is the final host that the attacker accesses and at what time? Choose the host and enter the time below (HH:MM:SS AM/PM).

12. From which host and at what time does the final piece of the attack chain originate? Choose the host and enter the time below (HH:MM:SS AM/PM).

13. Using the network data provided to you, determine the third suspicious remote host IP address that communicates with an internal host and at what time. Enter the IP address and time below (HH:MM:SS AM/PM).

14. What internal hosts appear to be compromised within the network? Check all that apply.

15. Describe in your own words the attack chain for the incident that occurred on this network. Please be as specific as possible by using hostnames, timestamps, and IP addresses.

16. What action would you take at this point?

B.1.3 Qualitative Feedback

I used a 5-point Likert scale to collect feedback for three different categories with the possible responses being Strongly Disagree, Somewhat Disagree, Neutral, Somewhat Agree, or Strongly Agree.

1. Ease of use:

(a) It was easy to maneuver through time.

(b) It was easy to see the traffic between hosts.

(c) It was easy to change node properties (colors, labels, or shapes).

(d) It was easy to select a node(s).

(e) It was easy to get more information about a host.

(f) It was easy get more information about a piece of network traffic.

(g) It was hard to use the tool.

2. Experience:

(a) I enjoyed using the visualization tool.

(b) I disliked the visualization layout.

(c) I could see myself using the tool in a professional environment.

(d) I could see myself using the tool for personal means.

(e) I would recommend this tool to someone else.

3. Efficiency:

(a) I experienced performance issues with the tool.

(b) I completed all of the required tasks easily with the tool provided.

(c) I completed some tasks with ease but struggled with a few.

(d) I struggled to complete any of the tasks with the tool provided.

Additionally, I allowed participants to type into a free-form text box or verbally provide feedback at the conclusion of their session.

B.2 Riverside Instructions

I allowed users to download the instructions located below to reference during the scenario, which supplemented the instructional video they watched before starting the scenario tasks.

This is a list of interactive options that you have as a user of the visualization tool as shown in the training video. Feel free to reference this as you complete tasks and answer questions.

Below is an overview of Riverside with **agents in the cyan color** and **remote nodes in the color gray**, connected by edges which represent different types of network traffic.

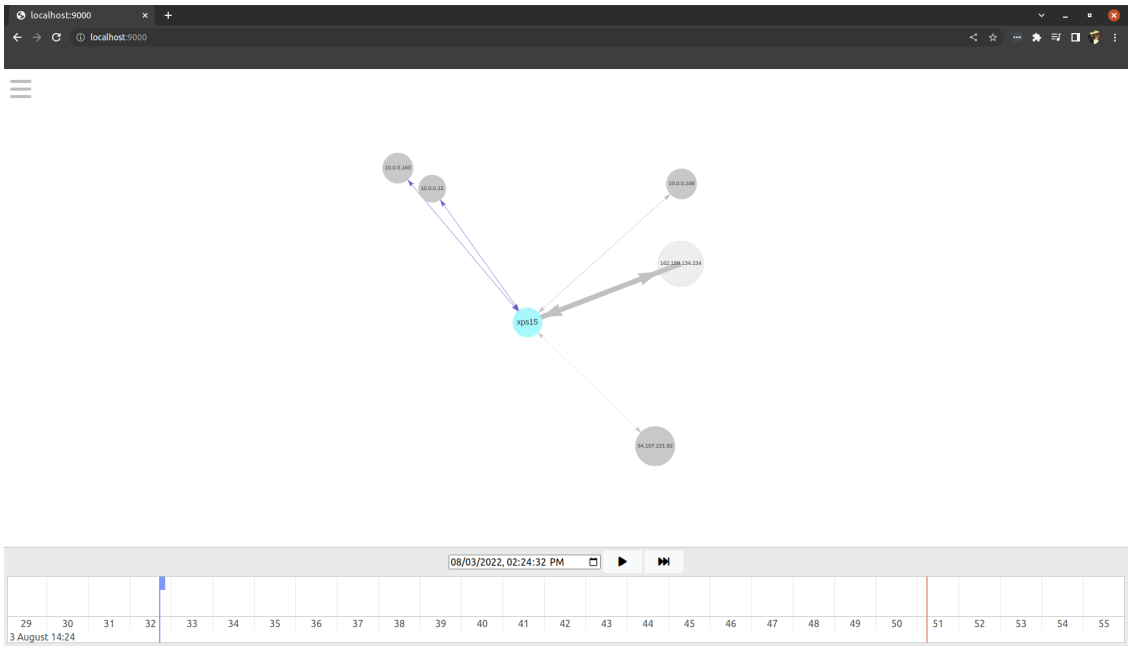

1. **Nodes can be dragged around on the visualization, which will move the edge connecting nodes as well. The visualization canvas can also be moved around and zoomed out or in to change what's in your canvas view.**
2. **<u>Hover:</u>** You can hover over a segment of network traffic or node to see more information about that component.

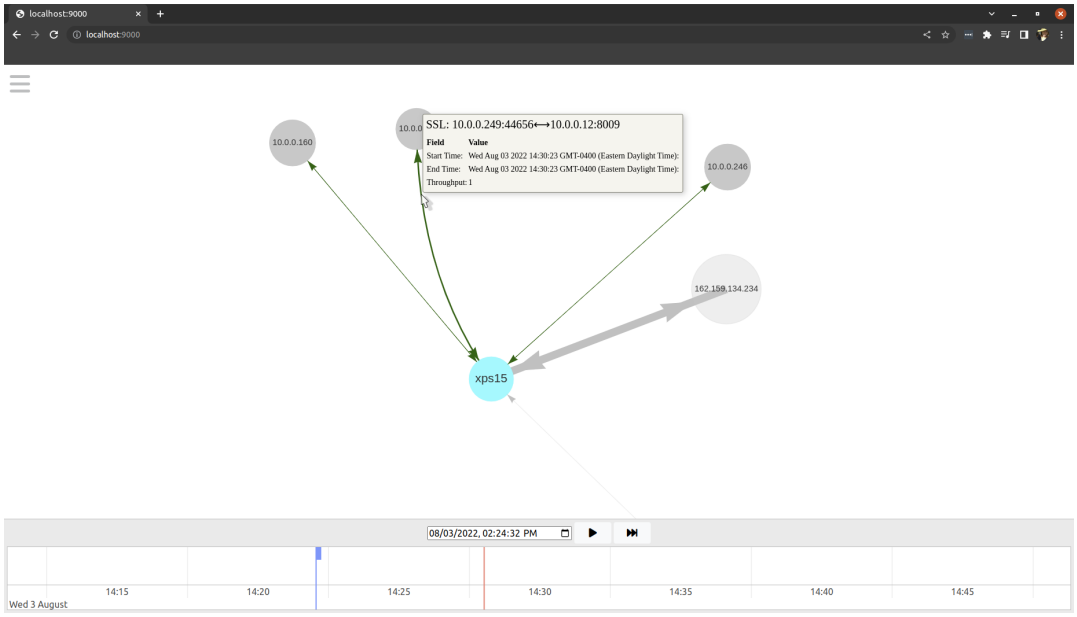

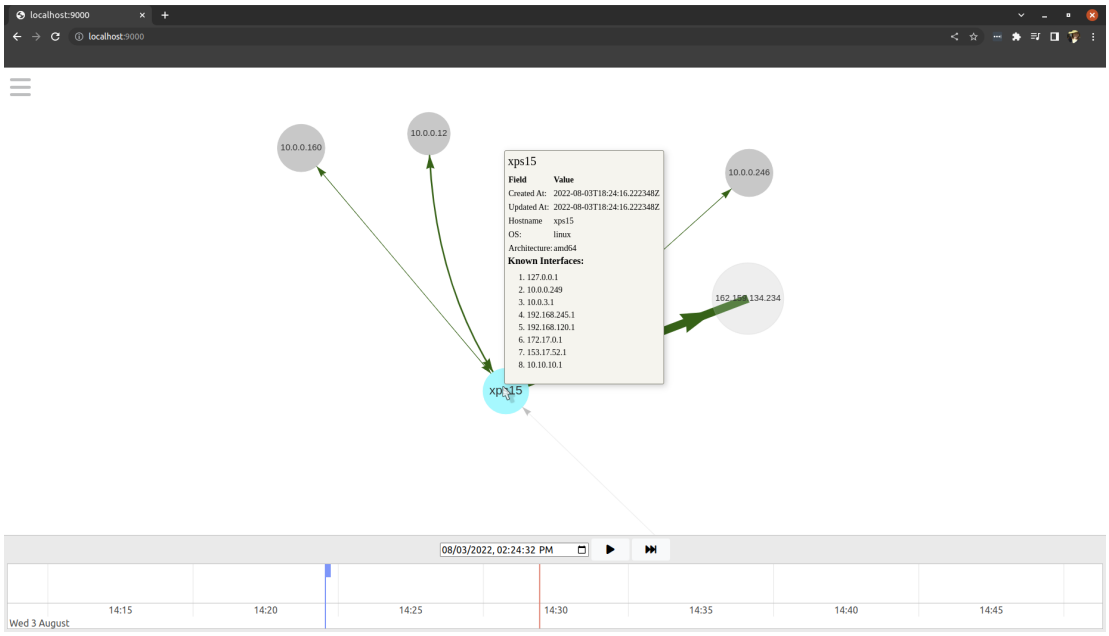

3. **Nodes:** You can change node colors, shapes, and labels by right-clicking it and choosing an option.
   a. **Colors:**

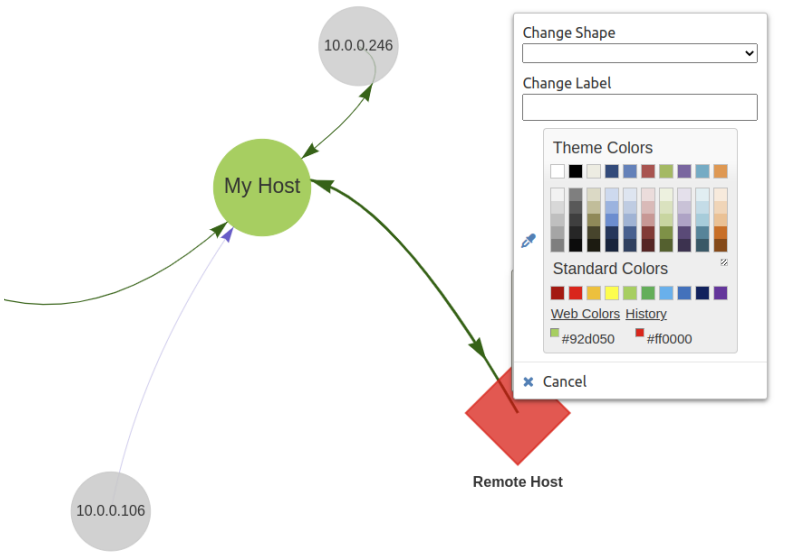

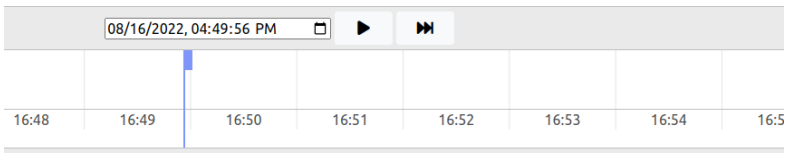

   b. **Labels:**

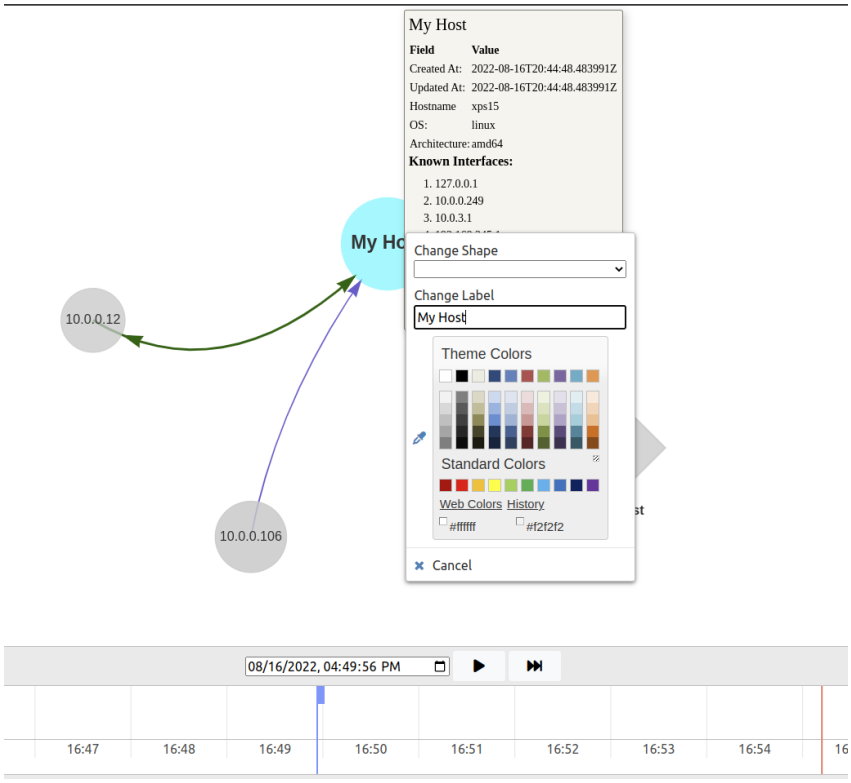
My Host
Field Value
Created At: 2022-08-16T20:44:48.483991Z
Updated At: 2022-08-16T20:44:48.483991Z
Hostname xps15
OS: linux
Architecture: amd64
Known Interfaces:
1. 127.0.0.1
2. 10.0.0.249
3. 10.0.3.1
My Ho
Change Shape
Change Label
My Host
Theme Colors
Standard Colors
Web Colors History
#ffffff
#f2f2f2
Cancel
10.0.0.12
10.0.0.106
08/16/2022, 04:49:56 PM
16:47
16:48
16:49
16:50
16:51
16:52
16:53
16:54

c. Shapes:

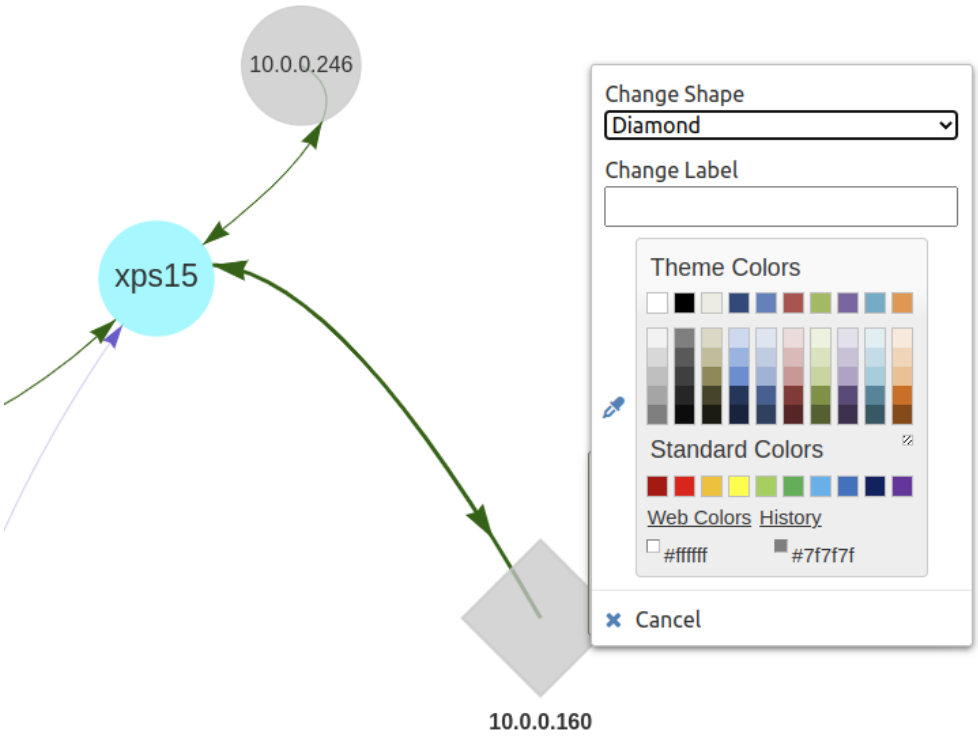
10.0.0.246
xps15
10.0.0.160
Change Shape
Diamond
Change Label
Theme Colors
Standard Colors
Web Colors History
#ffffff
#7f7f7f
Cancel

4. **Timeline:**
    a. You can scale the timeline by zooming in or out when hovering over it.

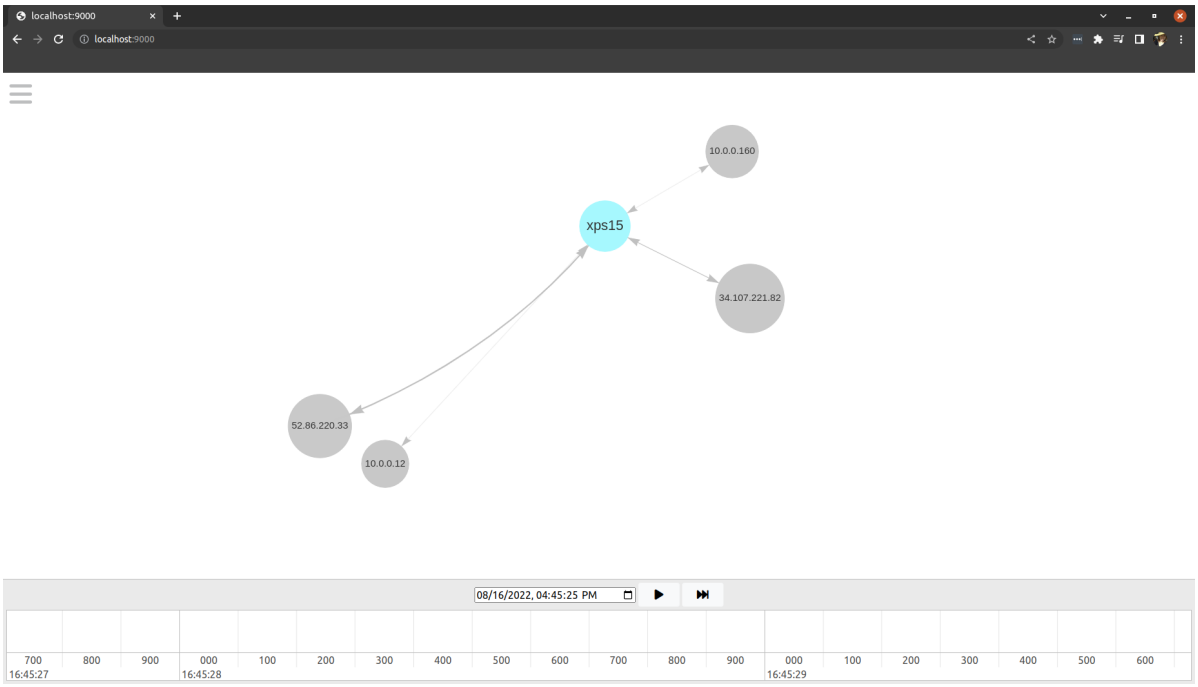

    b. The red line represents real, or current, time, where the blue cursor can be dragged when the visualization is paused to move forward or backward in time. Additionally you can enter a date and time in the box just above the timeline to automatically move to a point in time.

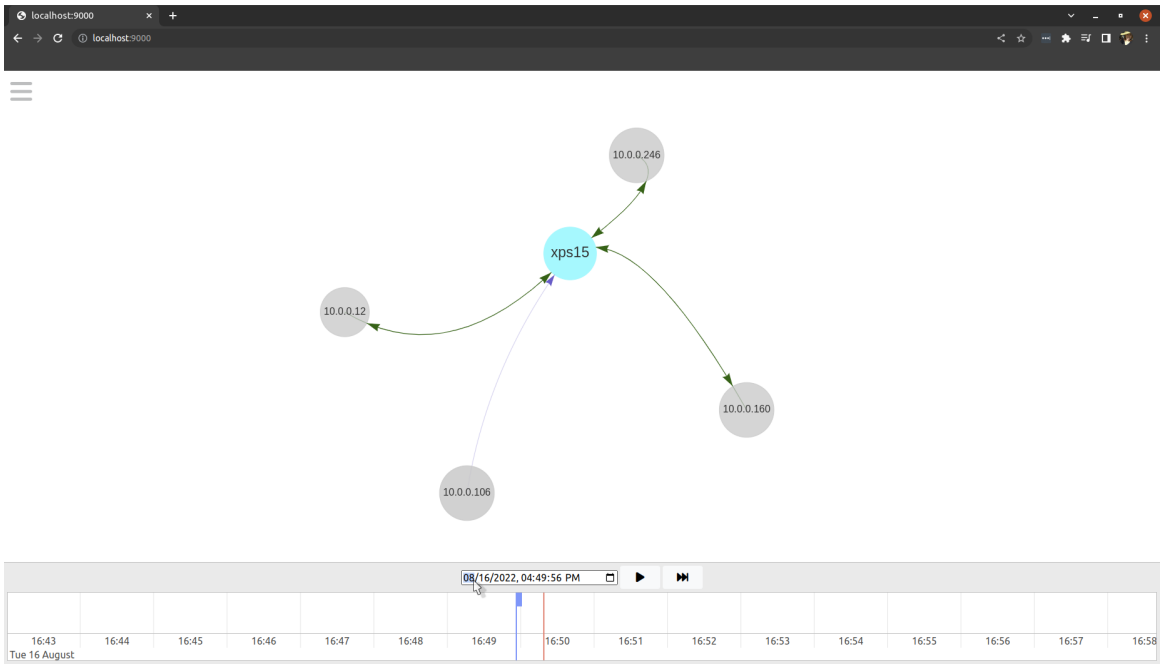

Bibliography

[1] M. W. Boyce, K. M. Duma, L. J. Hettinger, T. B. Malone, D. P. Wilson, and J. Lockett-Reynolds. Human Performance in Cybersecurity: A Research Agenda. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 55:1115–1119, Sep 2011.

[2] Faris Bugra Kokulu, Ananta Soneji, Tiffany Bao, Yan Shoshitaishvili, Ziming Zhao, Adam Doupé, and Gail-Joon Ahn. Matched and Mismatched SOCs: A Qualitative Study on Security Operations Center Issues. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 1955–1970. ACM, Nov 2019.

[3] Lyndsey Franklin, Meg Pirrung, Leslie Blaha, Michelle Dowling, and Mi Feng. Toward a Visualization-Supported Workflow for Cyber Alert Management Using Threat Models and Human-Centered Design. In *Proceedings of the 2017 IEEE Symposium on Visualization for Cyber Security*, VizSec '17, pages 1–8. IEEE, Oct 2017.

[4] Mica R. Endsley. Toward a Theory of Situation Awareness in Dynamic Systems. *Human Factors*, 37(1):32–64, Mar 1995.

[5] Michael Tyworth, Nicklaus A. Giacobe, and Vincent Mancuso. Cyber Situation Awareness as Distributed Socio-Cognitive Work. In *Proceedings of SPIE Cyber Sensing 2012*, volume

8408, pages 84080F–1 – 84080F–9. International Society for Optics and Photonics, SPIE, May 2012.

[6] Ulrik Franke and Joel Brynielsson. Cyber situational awareness – A systematic review of the literature. *Computers & Security*, 46:18–31, Oct 2014.

[7] Cyril Onwubiko and Thomas Owens. Review of Situational Awareness for Computer Network Defense. In *Situational Awareness in Computer Network Defense: Principles, Methods and Applications*, pages 1–9. IGI Global, 2012.

[8] Chaomei Chen. Information Visualization. *Wiley Interdisciplinary Reviews: Computational Statistics*, 2(4):387–403, Jul 2010.

[9] A. D'Amico and M. Kocka. Information Assurance Visualizations for Specific Stages of Situational Awareness and Intended Uses: Lessons Learned. In *Proceedings of the 2005 IEEE Workshop on Visualization for Computer Security*, VizSec '05, pages 107–112. IEEE, Oct 2005.

[10] Marc Grégoire and Luc Beaudoin. Visualization for Network Situational Awareness in Computer Network Defence. In *Proceedings of the Visualization and the Common Operational Picture*, pages 20–1–20–6, Jun 2005.

[11] Valérie Lavigne and Denis Gouin. Applicability of Visual Analytics to Defence and Security Operations. Technical report, Defense Research and Development Canada Valcartier (QUEBEC), 2011.

[12] Michael Tyworth, Nicklaus A. Giacobe, Vincent Mancuso, and Christopher Dancy. The Distributed Nature of Cyber Situation Awareness. In *Proceedings of the 2012 IEEE Inter-*

national Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support, page 174–178. IEEE, Mar 2012.

[13] R. Jordan Crouser, Erina Fukuda, and Subashini Sridhar. Retrospective on a Decade of Research in Visualization for Cybersecurity. In *Proceedings of the 2017 IEEE International Symposium on Technologies for Homeland Security*, page 1–5. IEEE, Apr 2017.

[14] Sean McKenna, Diane Staheli, and Miriah Meyer. Unlocking User-Centered Design Methods for Building Cyber Security Visualizations. In *Proceedings of the 2015 IEEE Symposium on Visualization for Cyber Security*, VizSec '15, page 1–8. IEEE, Oct 2015.

[15] Glenn A. Fink, Christopher L. North, Alex Endert, and Stuart Rose. Visualizing Cyber Security: Usable Workspaces. In *Proceedings of the 6th International Workshop on Visualization for Cyber Security*, page 45–56. IEEE, Oct 2009.

[16] Vinícius Tavares Guimarães, Carla Maria Dal Sasso Freitas, Ramin Sadre, Liane Margarida Rockenbach Tarouco, and Lisandro Zambenedetti Granville. A Survey on Information Visualization for Network and Service Management. *IEEE Communications Surveys Tutorials*, 18(1):285–323, Jul 2016.

[17] Diane Staheli, Tamara Yu, R. Jordan Crouser, Suresh Damodaran, Kevin Nam, David O'Gwynn, Sean McKenna, and Lane Harrison. Visualization Evaluation for Cyber Security: Trends and Future Directions. In *Proceedings of the 11th Workshop on Visualization for Cyber Security*, VizSec '14, page 49–56. ACM, Nov 2014.

[18] Anita D'Amico, Laurin Buchanan, Drew Kirkpatrick, and Paul Walczak. Cyber Operator Perspectives on Security Visualization. In *Advances in Human Factors in Cybersecurity*, volume 501, pages 69–81. Springer International Publishing, Jul 2016.

[19] NIST: Computer Security Resource Center Glossary. `https://csrc.nist.gov/glossary/term/defense_in_breadth`.

[20] Daniel M. Best, Alex Endert, and Daniel Kidwell. 7 Key Challenges for Visualization in Cyber Network Defense. In *Proceedings of the 11th Workshop on Visualization for Cyber Security*, VizSec '14, page 33–40. ACM, Nov 2014.

[21] Dustin L. Arendt, Russ Burtner, Daniel M. Best, Nathan D. Bos, John R. Gersh, Christine D. Piatko, and Celeste Lyn Paul. Ocelot: User-Centered Design of a Decision Support Visualization for Network Quarantine. In *Proceedings of the 2015 IEEE Symposium on Visualization for Cyber Security*, VizSec '15, page 1–8. IEEE, Oct 2015.

[22] Sneha Gathani, Shayan Monadjemi, Alvitta Ottley, and Leilani Battle. A Grammar-Based Approach for Applying Visualization Taxonomies to Interaction Logs. *Computer Graphics Forum*, 41(3):489–500, Jul 2022.

[23] Fabian Beck, Michael Burch, Stephan Diehl, and Daniel Weiskopf. A Taxonomy and Survey of Dynamic Graph Visualization. *Computer Graphics Forum*, 36(1):133–159, 2016.

[24] Jeff Sauro and James R Lewis. *Quantifying the User Experience: Practical Statistics for User Research*. Elsevier Science & Technology, Dec 2016.

[25] Carla M. D. S. Freitas, Marcelo S. Pimenta, and Dominique L. Scapin. User-Centered Evaluation of Information Visualization Techniques: Making the HCI-InfoVis Connection

Explicit. In *Handbook of Human Centric Visualization*, page 315–336. Springer New York, 2014.

[26] Kaitlyn DeValk and Niklas Elmqvist. Riverside: Dynamic Visualization of Network Traffic for Situation Awareness in Computer Security. In *Proceedings of 2022 IEEE Symposium on Visualization for Cyber Security*, Oct 2022. Poster presentation at 2022 VizSec Conference.

[27] Kaitlyn DeValk. Riverside: A Network Security Visualization Tool. Toolshed Presentation at 2022 Wild West Hackin' Fest Conference., Oct 2022.

[28] Simson Garfinkel and Heather Richter Lipford. *Usable Security: History, Themes, and Challenges*. Morgan & Claypool, Sept 2014.

[29] Christian Tiefenau, Maximilian Häring, Katharina Krombholz, and Emanuel von Zezschwitz. Security, Availability, and Multiple Information Sources: Exploring Update Behavior of System Administrators. In *Proceedings of the 16th Symposium on Usable Privacy and Security*, pages 239–258. USENIX Association, Aug 2020.

[30] Noura Alomar, Edward Qiu, Primal Wijesekera, Amit Elazari, and Serge Egelman. Poster: Bug Bounty Hunter, Red Teamer, or Pen tester? A Closer Look at the Roles of Security Teams in Vulnerability Discovery. In *Proceedings of the 15th Symposium on Usable Privacy and Security*, Feb 2019. Poster presented at the 15th Symposium on Usable Privacy and Security.

[31] John R. Goodall, Wayne G. Lutters, and Anita Komlodi. I Know My Network: Collaboration and Expertise in Intrusion Detection. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work*, pages 342–345. ACM, Nov 2004.

[32] Celeste Lyn Paul. Human-Centered Study of a Network Operations Center: Experience Report and Lessons Learned. In *Proceedings of the 2014 ACM Workshop on Security Information Workers*, WSIW '14, pages 39–42. ACM, Nov 2014.

[33] Cleidson R. B. de Souza, Claudio S. Pinhanez, and Victor F. Cavalcante. Information Needs of System Administrators in Information Technology Service Factories. In *Proceedings of the 5th ACM Symposium on Computer Human Interaction for Management of Information Technology*, page 1–10. ACM, Nov 2011.

[34] David Botta, Rodrigo Werlinger, André Gagné, Konstantin Beznosov, Lee Iverson, Sidney Fels, and Brian Fisher. Towards Understanding IT Security Professionals and Their Tools. In *Proceedings of the 3rd Symposium on Usable Privacy and Security*, pages 100–111. ACM, Jul 2007.

[35] Celeste Lyn Paul and Kirsten Whitley. A Taxonomy of Cyber Awareness Questions for the User-Centered Design of Cyber Situation Awareness. In *Human Aspects of Information Security, Privacy, and Trust*, pages 145–154. Springer, Berlin, Heidelberg, 2013.

[36] Anita D'Amico, Kirsten Whitley, Daniel Tesone, Brianne O'Brien, and Emilie Roth. Achieving Cyber Defense Situational Awareness: A Cognitive Task Analysis of Information Assurance Analysts. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 49(3):229–233, Sept 2005.

[37] Robert F. Erbacher, Deborah A. Frincke, Pak Chung Wong, Sarah Moody, and Glenn Fink. A Multi-Phase Network Situational Awareness Cognitive Task Analysis. *Information Visualization*, 9(3):204–219, Sep 2010.

[38] A. D'Amico and K. Whitley. The real work of computer network defense analysts. In *Proceedings of the Workshop on Visualization for Computer Security*, pages 19–37. Springer, Berlin, Heidelberg, 2008.

[39] Hadi Shiravi, Ali Shiravi, and Ali A. Ghorbani. A Survey of Visualization Systems for Network Security. *IEEE Transactions on Visualization and Computer Graphics*, 18(8):1313–1329, Aug 2012.

[40] Lane Harrison and Aidong Lu. The Future of Security Visualization: Lessons from Network Visualization. *IEEE Network*, 26:6–11, Nov 2012.

[41] Guanyao Du, Chun Long, Jianjun Yu, Wei Wan, Jing Zhao, and Jinxia Wei. A Real-time Big Data Framework for Network Security Situation Monitoring. In *Proceedings of the 21st International Conference on Enterprise Information Systems*, page 167–175. Science and Technology Publications, May 2019.

[42] Alex Ulmer, David Sessler, and Jörn Kohlhammer. NetCapVis: Web-based Progressive Visual Analytics for Network Packet Captures. In *Proceedings of the 2019 IEEE Symposium on Visualization for Cyber Security*, VizSec '19, page 1–10. IEEE, Oct 2019.

[43] Juraj Uhlár, Martin Holkovič, and Vít Rusňák. PCAPFunnel: A Tool for Rapid Exploration of Packet Capture Files. In *25th International Conference Information Visualization*, page 69–76. IEEE, Jul 2021.

[44] C.P. Lee, J. Trost, N. Gibbs, Raheem Beyah, and J.A. Copeland. Visual Firewall: Real-Time Network Security Monitor. In *Proceedings of the 2005 IEEE Workshop on Visualization for Computer Security*, VizSec '05, page 129–136. IEEE, Oct 2005.

[45] Trenton Bond. Visualizing Firewall Log Data to Detect Security Incidents. *Global Information Assurance Certification Paper*, pages 722–729, 2009.

[46] D. Phan, J. Gerth, M. Lee, A. Paepcke, and T. Winograd. Visual Analysis of Network Flow Data with Timelines and Event Plots. In *Proceedings of the Workshop on Visualization for Computer Security*, page 85–99. Springer, Berlin, Heidelberg, 2008.

[47] John R. Goodall, Eric D. Ragan, Chad A. Steed, Joel W. Reed, G. David Richardson, Kelly M.T. Huffer, Robert A. Bridges, and Jason A. Laska. Situ: Identifying and Explaining Suspicious Behavior in Networks. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):204–214, Jan 2019.

[48] Danny Yuxing Huang, Noah Apthorpe, Frank Li, Gunes Acar, and Nick Feamster. IoT Inspector: Crowdsourcing Labeled Network Traffic from Smart Home Devices at Scale. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 4(2):46:1–46:21, Jun 2020.

[49] Igor Kotenko and Evgenia Novikova. Vissecanalyzer: A visual analytics tool for network security assessment. In *Security Engineering and Intelligence Informatics*, volume 8128 of *Lecture Notes in Computer Science*, page 345–360. Springer, Berlin, Heidelberg, 2013.

[50] Frederik L Dennig, Eren Cakmak, Henrik Plate, and Daniel A Keim. VulnEx: Exploring Open-Source Software Vulnerabilities in Large Development Organizations to Understand

Risk Exposure. In *Proceedings of the 2021 IEEE Symposium on Visualization for Cyber Security*, VizSec '21, pages 79–83. IEEE, Oct 2021.

[51] Marco Angelini, Graziano Blasilli, Silvia Bonomi, Simone Lenti, Alessia Palleschi, Giuseppe Santucci, and Emiliano De Paoli. BUCEPHALUS: a BUsiness CEntric cybersecurity Platform for proActive anaLysis Using visual analyticS. In *Proceedings of the 2021 IEEE Symposium on Visualization for Cyber Security*, VizSec '21, page 15–25. IEEE, Oct 2021.

[52] Sadie Creese, Michael Goldsmith, Nick Moffat, Jassim Happa, and Ioannis Agrafiotis. CyberVis: Visualizing the Potential Impact of Cyber Attacks on the Wider Enterprise. In *2013 IEEE International Conference on Technologies for Homeland Security*, page 73–79. IEEE, Nov 2013.

[53] Jonathan McPherson, Kwan-Liu Ma, Paul Krystosk, Tony Bartoletti, and Marvin Christensen. PortVis: A Tool for Port-Based Detection of Security Events. In *Proceedings of the 2004 ACM Workshop on Visualization and Data Mining for Computer Security*, VizSEC/DMSEC '04, page 73–81. ACM, Oct 2004.

[54] Christopher Kintzel, Johannes Fuchs, and Florian Mansmann. Monitoring Large IP Spaces with ClockView. In *Proceedings of the 8th International Symposium on Visualization for Cyber Security*, VizSec '08, page 1–10. ACM, Jul 2011.

[55] Kiran Lakkaraju, Ratna Bearavolu, A. Slagell, W. Yurcik, and S. North. Closing-the-Loop in NVisionIP: Integrating Discovery and Search in Security Visualizations. In *Proceedings*

of the 2005 IEEE Workshop on Visualization for Computer Security, VizSec '05, page 75–82. IEEE, Oct 2005.

[56] Florian Mansmann, Fabian Fischer, Daniel A. Keim, Stephan Pietzko, and Marcel Waldvogel. Interactive Analysis of NetFlows for Misuse Detection in Large IP Networks. In *DFN-Forum Kommunikationstechnologie*, Garching (Munich), Germany, 2009.

[57] Troy Nunnally, Kulsoom Abdullah, A. Selcuk Uluagac, John A. Copeland, and Raheem Beyah. NAVSEC: A Recommender System for 3D Network Security Visualizations. In *Proceedings of the 10th Workshop on Visualization for Cyber Security*, page 41–48. ACM, Oct 2013.

[58] Daisuke Inoue, Masashi Eto, Koei Suzuki, Mio Suzuki, and Koji Nakao. DAEDALUS-VIZ: Novel Real-Time 3D Visualization for Darknet Monitoring-based Alert System. In *Proceedings of the 9th International Symposium on Visualization for Cyber Security*, page 72–79. ACM, Oct 2012.

[59] Wireshark. `https://www.wireshark.org/`.

[60] Fangfang Zhou, Wei Huang, Ying Zhao, Yang Shi, Xing Liang, and Xiaoping Fan. ENTVis: A Visual Analytic Tool for Entropy-Based Network Traffic Anomaly Detection. *IEEE Computer Graphics and Applications*, 35(6):42–50, Nov 2015.

[61] J.R. Goodall, W.G. Lutters, P. Rheingans, and A. Komlodi. Focusing on Context in Network Traffic Analysis. *IEEE Computer Graphics and Applications*, 26(2):72–80, Mar 2006.

[62] Kiran Lakkaraju, William Yurcik, and Adam J. Lee. NVisionIP: NetFlow Visualizations of System State for Security Situational Awareness. In *Proceedings of the 2004 ACM Workshop on Visualization and Data Mining for Computer Security*, VizSEC/DMSEC '04, page 65–72. ACM, Oct 2004.

[63] Xiaoxin Yin, William Yurcik, Michael Treaster, Yifan Li, and Kiran Lakkaraju. VisFlowConnect: NetFlow Visualizations of Link Relationships for Security Situational Awareness. In *Proceedings of the 2004 ACM Workshop on Visualization and Data Mining for Computer Security*, VizSEC/DMSEC '04, pages 26–34. ACM, Oct 2004.

[64] Robert Ball, Glenn A. Fink, and Chris North. Home-Centric Visualization of Network Traffic for Security Administration. In *Proceedings of the 2004 ACM Workshop on Visualization and Data Mining for Computer Security*, VizSEC/DMSEC '04, page 55–64. ACM, Oct 2004.

[65] Teryl Taylor, Diana Paterson, Joel Glanfield, Carrie Gates, Stephen Brooks, and John McHugh. FloVis: Flow Visualization System. In *Proceedings of the 2009 Cybersecurity Applications and Technology Conference for Homeland Security*, page 186–198, Mar 2009.

[66] Hideki Koike, Kazuhiro Ohno, and Kanba Koizumi. Visualizing Cyber Attacks using IP Matrix. In *Proceedings of the 2005 IEEE Workshop on Visualization for Computer Security*, VizSec '05, page 91–98, Oct 2005.

[67] Joel Glanfield, Stephen Brooks, Teryl Taylor, Diana Paterson, Christopher Smith, Carrie Gates, and John McHugh. OverFlow: An Overview Visualization for Network Analysis.

In *6th International Workshop on Visualization for Cyber Security*, page 11–19, Oct 2009.

[68] S. Foresti, J. Agutter, Y. Livnat, S. Moon, and R. Erbacher. Visual Correlation of Network Alerts. *IEEE Computer Graphics and Applications*, 26(2):48–59, Mar 2006.

[69] Fangfang Zhou, Ronghua Shi, Ying Zhao, Yezi Huang, and Xing Liang. NetSecRadar: A Visualization System for Network Security Situational Awareness. In *International Symposium on Cyberspace Safety and Security*, page 403–416. Springer International Publishing, 2013.

[70] Daniel M. Best, Shawn Bohn, Douglas Love, Adam Wynne, and William A. Pike. Real-time Visualization of Network Behaviors for Situational Awareness. In *Proceedings of the 7th International Symposium on Visualization for Cyber Security*, VizSec '10, page 79–90. ACM, Sept 2010.

[71] Robert F. Erbacher. Visualization Design for Immediate High-Level Situational Assessment. In *Proceedings of the 9th International Symposium on Visualization for Cyber Security*, VizSec '12, page 17–24. ACM, Oct 2012.

[72] Elisha Peterson. Dagger: Modeling and Visualization for Mission Impact Situation Awareness. In *Proceedings of the 2016 IEEE Military Communications Conference*, page 25–30, Nov 2016.

[73] Fabian Fischer, Johannes Fuchs, Florian Mansmann, and Daniel A Keim. BANKSAFE: A Visual Situational Awareness Tool for Large-Scale Computer Networks. *Information Visualization*, 14(1):51–61, Jan 2015.

[74] Siming Chen, Cong Guo, Xiaoru Yuan, Fabian Merkle, Hanna Schaefer, and Thomas Ertl. OCEANS: Online Collaborative Explorative Analysis on Network Security. In *Proceedings of the 11th Workshop on Visualization for Cyber Security*, VizSec '14, page 1–8. ACM, Nov 2014.

[75] Celeste Lyn Paul, Randall Rohrer, Patrick Sponaugle, Jenna Huston, and Bohdan Nebesh. CyberSAVI: A Cyber Situation Awareness Visual Interface for Mission-Level Network Situation Awareness. In *Proceedings of the 2013 IEEE Symposium on Visualization for Cyber Security*, VizSec '13, Oct 2013.

[76] Fabian Fischer and Daniel A. Keim. NStreamAware: Real-Time Visual Analytics for Data Streams to Enhance Situational Awareness. In *Proceedings of the 11th Workshop on Visualization for Cyber Security*, page 65–72. ACM, Nov 2014.

[77] Dustin Arendt, Dan Best, Russ Burtner, and Celeste Lyn Paul. CyberPetri at CDX 2016: Real-time Network Situation Awareness. In *Proceedings of the 2016 IEEE Symposium on Visualization for Cyber Security*, VizSec '16. IEEE, Oct 2016.

[78] Cameron C. Gray, Panagiotis D. Ritsos, and Jonathan C. Roberts. Contextual Network Navigation to Provide Situational Awareness for Network Administrators. In *Proceedings of the 2015 IEEE Symposium on Visualization for Cyber Security*, VizSec '15, Oct 2015.

[79] Gephi: The Open Graph Viz Platform. `https://gephi.org/`.

[80] Cytoscape. `https://cytoscape.org/`.

[81] Splunk Dashboards and Visualizations. `https://docs.splunk.com/Documentation/Splunk/latest/Viz/Visualizationreference`.

[82] Arkime. https://github.com/arkime/arkime.

[83] What is Kibana? https://www.elastic.co/what-is/kibana.

[84] VizAlerts: Data-driven alerting for Tableau Server. https://github.com/tableau/VizAlerts, Oct 2015.

[85] Vincent J Ercolani, Mark W Patton, and Hsinchun Chen. Shodan Visualized. In *2016 IEEE Conference on Intelligence and Security Informatics*, page 193–195, Sep 2016.

[86] Skydive-Project. Skydive-project/skydive: An open source real-time network topology and protocols analyzer. https://github.com/skydive-project/skydive, Feb 2016.

[87] Jim Accord. Situational Awareness and ICS using GRASSMARLIN, Nov 2017.

[88] Aneesha Sethi, Federica Paci, and Gary Wills. EEVi - Framework for Evaluating the Effectiveness of Visualization in Cyber-Security. *International Journal of Intelligent Computing Research (IJICR)*, 7(4):761–770, Dec 2016. The original source for the publication is the "International Journal of Intelligent Computing Research" and the publisher is "Infonomics Society.".

[89] Aneesha Sethi and Gary Wills. Expert-Interviews Led analysis of EEVi — A Model for Effective Visualization in Cyber-Security. In *Proceedings of the 2017 IEEE Symposium on Visualization for Cyber Security*, VizSec '17. IEEE, Oct 2017.

[90] Iman Sharafaldin, Arash Habibi Lashkari, and Ali A. Ghorbani. An evaluation framework for network security visualizations. *Computers & Security*, 84:70–92, Jul 2019.

[91] S. McKenna, D. Staheli, C. Fulcher, and M. Meyer. BubbleNet: A Cyber Security Dashboard for Visualizing Patterns. *Computer Graphics Forum*, 35(3):281–290, Jun 2016.

[92] J. Stoll, D. McColgin, M. Gregory, V. Crow, and W. K. Edwards. Adapting Personas for Use in Security Visualization Design. In *VizSEC 2007: Proceedings of the Workshop on Visualization for Computer Security*, pages 39–52. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.

[93] Markus Wagner, Wolfgang Aigner, Alexander Rind, Hermann Dornhackl, Konstantin Kadletz, Robert Luh, and Paul Tavolato. Problem Characterization and Abstraction for Visual Analytics in Behavior-Based Malware Pattern Analysis. In *Proceedings of the 11th Workshop on Visualization for Cyber Security*, VizSec '14, page 9–16. ACM, Nov 2014.

[94] Silvia Miksch and Wolfgang Aigner. A matter of time: Applying a data–users–tasks design triangle to visual analytics of time-oriented data. *Computers & Graphics*, 38:286–290, Feb 2014.

[95] Philip A. Legg. Enhancing Cyber Situation Awareness for Non-Expert Users using Visual Analytics. In *2016 International Conference On Cyber Situational Awareness, Data Analytics and Assessment*, Jun 2016.

[96] Jeevitha Mahendiran, Kirstie A. Hawkey, and Nur Zincir-Heywood. Exploring the Need for Visualizations in System Administration Tools. In *CHI '14 Extended Abstracts on Human Factors in Computing Systems*, page 1429–1434. ACM, Apr 2014.

[97] Andrew Rinaldi. How much should your SMB budget for cybersecurity? https://www.business.com/articles/smb-budget-for-cybersecurity/, Sep 2022.

[98] Virginia Braun and Victoria Clarke. Using Thematic Analysis in Psychology. *Qualitative Research in Psychology*, 3:77–101, Jan 2006.

[99] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. Reliability and Inter-Rater Reliability in Qualitative Research: Norms and Guidelines for CSCW and HCI Practice. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW), Nov 2019.

[100] Johnny Saldaña. *The coding manual for qualitative researchers*. SAGE Publications, Los Angeles, California, 2nd edition, 2013.

[101] Ben Shneiderman. The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations. In *The Craft of Information Visualization*, Interactive Technologies, pages 364–371. Morgan Kaufmann, San Francisco, 2003.

[102] Purvi Saraiya, P. Lee, and C. North. Visualization of Graphs with Associated Timeseries Data. In *Proceedings of the 2005 IEEE Symposium on Information Visualization*, InfoVis '05, page 225–232, Oct 2005.

[103] Ilya Boyandin, Enrico Bertini, and Denis Lalanne. A Qualitative Study on the Exploration of Temporal Changes in Flow Maps with Animation and Small-Multiples. *Computer Graphics Forum*, 31(3):1005–1014, Jun 2012.

[104] Daniel J. Simons and Daniel T. Levin. Change Blindness. *Trends in Cognitive Sciences*, 1(7):261–267, 1997.

[105] Benjamin Bach, Emmanuel Pietriga, and Jean-Daniel Fekete. GraphDiaries: Animated Transitions and Temporal Navigation for Dynamic Networks. *IEEE Transactions on Visualization and Computer Graphics*, 20(5):740–754, May 2014.

[106] Eloïse Loubier and Bernard Dousset. Temporal and Relational Data Representation by Graph Morphing. *Safety and Reliability for Managing Risk*, 14(2), 2008.

[107] Tamara Munzner. *Visualization Analysis and Design*. CRC Press LLC, 2014.

[108] The WebSocket API. `https://developer.mozilla.org/en-US/docs/Web/API/WebSockets_API`.

[109] Google. Go: Build fast, reliable, and efficient software at scale. `https://go.dev/`.

[110] Introduction to gRPC. `https://grpc.io/docs/what-is-grpc/introduction/`, Aug 2021.

[111] Vis.js Community Edition. `https://visjs.org/`.

[112] Gorilla WebSocket. `https://github.com/gorilla/websocket`, May 2016.

[113] Gin Web Framework. `https://github.com/gin-gonic/gin`, May 2015.

[114] Jinzhu. Gorm. `https://gorm.io/`, 2013.

[115] MITRE ATT&CK. `https://attack.mitre.org/`.

[116] Cyber Security Analyst Demographics and Statistics: Number of Cybersecurity Analysts in the US. `https://www.zippia.com/cyber-security-analyst-jobs/demographics/`, Sep 2022.

[117] Arnaud Sallaberry, Chris Muelder, and Kwan-Liu Ma. Clustering, visualizing, and navigating for large dynamic graphs. In *Graph Drawing*, volume 7704 of *Lecture Notes in Computer Science*, page 487–498. Springer, Berlin, Heidelberg, 2013.

[118] Tarik Crnovrsanin, Shilpika, Senthil Chandrasegaran, and Kwan-Liu Ma. Staged Animation Strategies for Online Dynamic Networks. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):539–549, Feb 2021.

[119] Nicklaus A. Giacobe. A Picture is Worth a Thousand Alerts. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 57(1):172–176, Sep 2013.

[120] Jeff Sauro. Measuring Usability with the System Usability Scale (SUS), Feb 2011.