Cover-VT: Converged Security Visualization Tool

William Urbanski* Matthew Dunlop*† Randy Marchany† Joseph Tront*
*Bradley Department of Electrical and Computer Engineering
†Virginia Tech Information Technology Security Office
Virginia Tech, Blacksburg, VA 24061, USA
Email: {urbanski,dunlop,marchany,jgtront}@vt.edu

***Abstract*—The amount of data that floods today's networks is well beyond what security analysts can manage by textual means alone. In an effort to solve this problem, researchers have explored different methods of visualizing network security threats. There is little debate that humans can perceive more information visually than textually. The problem is that the majority of visualization tools in practice or proposed do not take efficient visualization techniques into consideration. As a result, it is difficult to get a high-level view of the network that facilitates rapid isolation of network attacks. We propose the Converged Security Visualization Tool (Cover-VT) to solve the efficient visualization problem. Cover-VT was designed to provide analysts with a high-level view of network threats using geographic information systems. The tool allows for rapid identification of threats by minimizing the cognitive obstacles to efficient threat location. Cover-VT includes the capability to drill-down on a node of interest for additional details and even filter out unwanted data. Cover-VT was designed with usability in mind, making it easy to comprehend while assisting the analyst in rapid threat identification. Many different security tools make up a security analyst's tool kit. Cover-VT was developed as an effective security visualization system that integrates existing security tools and network security systems.**

***Index Terms*—Converged Security, Security Management, Visualization**

I. Introduction

Information overload refers to a situation in which individuals are faced with more information than they can possibly process and absorb [12]. This is a common occurrence for security analysts who are bombarded with an enormous number of alerts every day. When individuals are faced with too much information they have to pick and choose which data to analyze, which often leads to important data getting overlooked [8].

Research has shown that humans can process much more information visually than textually [5], [14]. Consequently, research into security visualization has grown in popularity over recent years. In fact, there are a number of different approaches to tackling the issue of how to visualize security threats to computer networks. The goal of each of these visualization approaches is to present analysts with a picture that helps them better process large amounts of data. An effective visualization tool assists analysts in rapidly isolating specific threats without having to sacrifice data. The problem is that most tools are good at either rapidly isolating threats or rapidly identifying targets, but not both.

To address this issue, we propose the **Co**n**v**erged **Secu**r**i**ty **V**isualization **T**ool (Cover-VT). Our design leverages geographic information systems (GIS) to visualize hosts on the network while providing analysts with the capability to drill-down on specific threats or targets. By mapping nodes using GIS, analysts are provided with a familiar and predictable representation that aids in providing a visual understanding of network assets at a glance. Analysts can use the drill-down capability to see a more refined GIS picture, detailed graphs, or even the actual textual data itself. Analysts can also establish and save customizable filters to aid in monitoring key assets. All this is accomplished without losing any potentially important data.

The remainder of the paper is organized as follows. Section II compares Cover-VT with some of the more popular security visualization tools. This leads to Section III, where we detail the design of Cover-VT. In Section IV, we demonstrate some of the implementation details of Cover-VT. We then highlight some of the limitations of Cover-VT in Section V. In Sections VI and VII, we discuss future work and conclude, respectively.

II. Related Work

There are a myriad of visualization tools both in practice and in literature. We discuss some of the more popular tools and how they compare with Cover-VT. One such tool is the Open Source Security Information Management (OSSIM) [7]. OSSIM is essentially a collection of open source tools. OSSIM provides both textual-based analysis as well as visualization through the tools that it incorporates. Despite the large number of different charts and graphs, OSSIM does not provide a comprehensive snapshot of the network status. Cover-VT does this with the possibility of viewing every node in the network and its status using a scalable geospatial view. Another tool, SnortView [4], consists of three frames: the source address frame, alert frame, and source-destination matrix frame. The source address frame sequentially lists all sources for which there are alerts in the log. The alert frame presents the administrator with a time view of all alerts pertaining to a particular source address. Different alerts are visualized using different symbols and colors. The source-destination matrix frame provides a list of destinations in columns across the bottom of the frame. Cover-VT improves upon SnortView's visualization by displaying source addresses according to their geographic location instead of in sequential order. Cover-VT also improves upon SnortView's alert frame by minimizing the use of different features to indicate alerts.

978-1-4244-9221-3/11/$26.00 ©2011 IEEE

Two popular tools that attempt to visualize network anomalies rather than known signatures are RUMINT and AfterGlow. RUMINT [13] was designed to help analysts view and analyze a large number of network packets at one time. It includes seven different visualizations and can compare up to 19 different header fields simultaneously. Where Cover-VT improves upon RUMINT is in its ability to reduce clutter. Both the parallel coordinate plot and the binary rainfall visualization display a large amount of data at a time. They are good for seeing a shift in traffic patterns, but not good for rapidly pinpointing specific targets or malicious hosts. Cover-VT provides a global visualization that facilitates rapid identification of both target and potentially malicious hosts.

AfterGlow [1] can visualize network anomalies using either a linked graph or a treemap. A linked graph consists of machines or domains representing nodes connected by edges. Linked graphs do not produce predictable or intuitive visualizations. The problem is that linked graphs attempt to generate in such a way that edges do not cross. This means that nodes will likely show up in different locations every time. Cover-VT provides predictability by geographically positioning nodes. A treemap is similar to a tree structure. Children nodes are mapped inside of parent nodes and nodes at the same tree level are mapped adjacently. Treemaps become cluttered quickly when fed large amounts of varying data. On a large network, boxes can get so small that it is extremely difficult to tell what is going on. Cover-VT reduces clutter by clustering nodes in close proximity.

III. Cover-VT Design

Most security analysts today rely on a suite of tools to effectively assess network security threats. Examples of some of these tools are discussed in Section II. The problem with this approach is that the tools are disjointed, likely leaving gaps in analysis. To combat these difficulties Cover-VT was designed with usability in mind. In addition to designing an extremely intuitive interface, we incorporate multiple fine-grained analysis capabilities to alleviate the need for analysts to open up multiple, separate tools to make an assessment.

Modern network security infrastructures rely on layers of security devices to ensure complete network protection. Border firewalls may provide access control lists (ACLs) to block malicious hosts and known external threats while in-line Intrusion Prevention Systems (IPS) may filter content from network traffic in real time to protect end users. Cover-VT's subsystem has a modular design that allows it to accept data from multiple types of network security devices. Cover-VT integrates this data to generate a correlated report for each device. The correlated report contains a summary of the anomalous traffic detected by each security device. From within the report, links to the respective security devices' reporting systems are included in order to provide analysts with the capability to quickly generate extremely detailed reports from the applicable security system. By integrating multiple systems in this way, Cover-VT allows analysts to drill down beyond the visualization and utilize each security device's detailed reporting system. Cover-VT gives the analyst a holistic overview of anomalous traffic for a specific device without having to sacrifice details that can be helpful when analyzing a problem.

While many proprietary network security appliances require the use of bulky, operating system-dependent software to access their administration and reporting interfaces, Cover-VT has been implemented as a standards-compliant web application. Cover-VT relies on the Google Earth browser plugin and API as the primary geospatial interface. This is due to the intuitive navigation tools and high-resolution satellite imagery provided by the plugin. The Cover-VT application generates Keyhole Markup Language (KML) files for rendering in the browser plugin. Thus, Cover-VT is compatible with any Open Geospatial Consortium, Inc. (OGC) KML client [6].

A. GIS Benefits

As mentioned, Cover-VT is designed to integrate network security systems within a GIS interface. There are many benefits to using a representation that leverages GIS for security visualization. The primary advantage is that GIS representations are intuitive; people are familiar with looking at and reading maps. Therefore, a familiarity of the interface already exists.

One of the key capabilities built into Cover-VT is the merging of cyber security with physical security. This fusion is known as converged security. We have already mentioned some of the ways that Cover-VT assesses the cyber security of a network. Cover-VT can also monitor the physical security through the use of a geospatial representation. One way is by monitoring the status of co-located devices. A number of devices exhibiting anomalous behavior in the same geographic area could be an indication of a physical breach. Alternatively, the sudden appearance of an unexpected node near a critical system may be suspicious. By representing nodes according to their physical locations, analysts get a feel for what is happening at a specific geographic location. The physical security aspect of Cover-VT can even be enhanced with the capability to view camera displays at crucial locations or monitor the status of door locks.

Another advantage is that, since nodes are displayed according to their physical location, they show up in consistent, predictable locations. As a result, analysts get used to scanning certain areas of the map where critical systems are located. This is a significant improvement over a traditional text-based analysis system that requires analysts to recognize a critical system by its Internet Protocol (IP) address alone. On the Virginia Tech network (which is comprised of two /16 network blocks) identification of a critical server requires the analyst to recognize a significant IP address out of as many as 131,072 other addresses. As an alternative method, administrators often use intuitive names to help identify critical systems and their locations. Although these names may assist administrators, they also assist malicious users by telling them where the critical systems are located. Using GIS, administrators are able to monitor critical systems by their physical locations

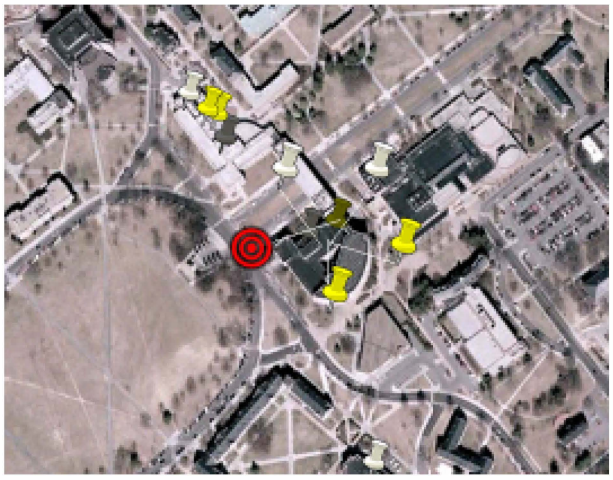

Fig. 1. Screenshot of a portion of the Cover-VT web-based application. The captured area displays the different alert levels.

rather than by potentially incriminating names. Since analysts will likely know the location of major server rooms, physical location can also help them infer the importance of targets. Machines that appear in a known server room are more likely to be closely scrutinized than machines of less value (e.g. in a dorm or in the library). Even with complex deployments that include clustering or the use of virtualization technologies, the ability to tie an IP address to its physical location greatly improves the quality of protection and assessment that analysts are able to provide to critical systems.

GIS also improves the ability to communicate and collaborate with others because analysts are better able to communicate the security posture of the network to non-technical audiences. With the intuitive display, analysts can quickly convey to management any network-related concerns that might have previously escaped understanding and interest.

B. *Efficient Visualization*

We took obstacles to preattentive processing into account when designing Cover-VT. Preattentive processing refers to the capability for a user to rapidly select targets by processing all the objects at once. One such obstacle is clutter, or objects masking each other. To prevent this, we implement what we call zoom clustering, which means at different elevations (zoom levels), objects in the same proximity will cluster into a single object [16]. This is particularly applicable in a tool like Cover-VT that uses GIS to maintain positional context. The observer is able to analyze a specific target by choosing a cluster and drilling-down to the target of interest [11].

Another technique we use that reduces clutter is to limit the number of identifying features between distractors as well as the number of distractors themselves. We use the term distractor to refer to any object that is not the target object. In Cover-VT, objects represent machines on, or attacking, the Virginia Tech network. These machines are categorized, by the network intrusion detection systems, into three groups based on severity of attack: low, medium, or high risk. Of these three, high risk machines are the ones that require immediate attention. To achieve this, we use three of the key features that researchers have identified to preserve preattentive detection: curvature, color, and size [2]. We use three features since research indicates that multiple unique features improve search capabilities [15]. To that end, we represent high risk machines with a red target symbol. The red target symbol, in addition to being a different color and shape, is slightly larger than the pushpin symbols used for distractor objects. Fig. 1 demonstrates how the red targets stand-out on the GIS display, thus achieving preattentive detection capability.

We also designed Cover-VT to mitigate the effects of change blindness [9]. We were concerned that when a machine changed from low or medium risk to high, that an analyst might not notice. This is especially true if the analyst is not actively observing the display. To account for this, we temporarily apply two additional features to the high priority target. The first feature is again size. We have the target symbol of new high risk targets appear much larger than stale targets. As the new target transitions to stale, the size will decay to that of other stale targets. The second feature we incorporate is an audible alarm. This is to redirect the analyst's attention if it is not focused on the display. It also serves to add a multimodal aspect to Cover-VT in the event that an analyst's cognitive threshold is being stressed [3].

IV. Implementation

Cover-VT is being used as an operational security tool in the Virginia Tech Information Technology Security Lab. It is the primary traffic classification and reporting system for a campus-wide Intrusion Detection System (IDS) based on Snort [10] and other open-source tools. Using Global Positioning System (GPS) data provided by the Virginia Tech Enterprise GIS Research and Development Administration we are able to determine the physical location of each room on campus. These locations, combined with a live database of network device registrations, enable the mapping of each computer to a room on the Virginia Tech campus. Network attack data generated by the IDS sensors is merged with the GIS repository to allow us to represent network attacks on campus in a geospatial interface. Fig. 1 shows a screenshot of the working Cover-VT prototype web-based application. We are also working to integrate 3D building models and floor plans as shown in Fig. 2.

Cover-VT provides global coverage of the Virginia Tech campus while still having the capability for fine-grained analysis and customization. Our tool identifies all internal and external network attacks because of sensors located on core network routers as well as on the network border. These sensors are also able to detect and visualize attacks that occur between machines on campus. Cover-VT's modular design allows us to integrate our existing security tools into the framework. Threats can be weighted by the type of system and the Snort alert classification to determine what constitutes high threat alerts. Therefore, critical systems will trigger high threat alerts before non-critical systems. This weighting function

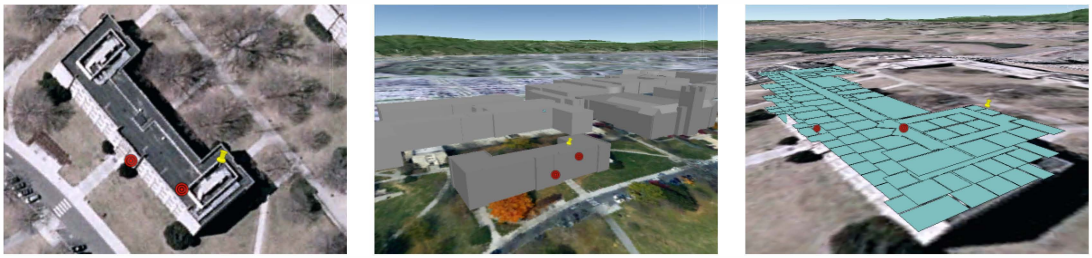

(a) Overhead satellite view. (b) 3D structure view. (c) 3D floor break-out view.

Fig. 2. Different zoom levels within Cover-VT. Fig. 2(a) depicts the standard overhead view as a user zooms in on a structure. The next level of detail is shown in Fig. 2(b) where 3D structures are represented. In Fig. 2(c), the analyst has the ability to view the physical location of the alerts down to the actual floor and room. (*Images © 2010 Google, Commonwealth of Virginia, Digital Globe, USDA Farm Service Agency*)

can easily be tailored to the needs of the analyst. Through the use of callbacks we have the ability to integrate with existing trouble ticket systems, external databases, and other applications currently being utilized as a part of the security notification process.

V. Limitations

Devices that perform Network Address Translation (NAT) are troublesome for Cover-VT because they obscure the original location of the host. In these cases, attacks in Cover-VT are plotted as coming from the gateway device that is performing the NAT. This is not a problem limited to Cover-VT. External security analysts also cannot discern which host is performing a malicious action from behind a NAT.

Wireless users present a different set of challenges for Cover-VT because they have the ability to roam while they are transmitting to the network. Wireless users can freely associate and disassociate with wireless access points as they move around campus. Machines connected to a particular wireless access point (WAP) appear as co-located with that WAP.

VI. Future Work

A future enhancement that is currently being incorporating into Cover-VT is the ability to drill down from an overhead satellite view (Fig. 2(a)) to a 3D structure view (Fig. 2(b)) and then again to a floor-plan break-out view (Fig. 2(c)). These three different zoom levels provide analysts with the capability to visualize alerts within a particular structure even down to the individual room locations of the equipment. By providing this level of detail to the analyst, we enhance the physical security aspect of Cover-VT.

VII. Conclusion

The massive amount of information facing security analysts will continue to grow as networks become faster and as more network-capable devices are produced. Cover-VT was designed specifically to leverage cognitive principles that aid in rapid target selection. By doing this, we assist analysts in monitoring critical network threats. Our incorporation of GIS, produces a familiar and intuitive visual interface for analysts. This interface provides a broad overview of network assets with the capability to rapidly hone in on problem areas. Our telescoping design is combined with one-click filtering and different signature and anomaly based correlation engines to rapidly isolate specific network issues.

Using geospatial visualization for the management of security threats is a relatively unexplored research area. Cover-VT demonstrates how powerful a GIS front-end can be incorporated with fine-grained correlation engines to produce a single tool that meets all of a security analyst's needs. We propose that our tool will enhance the speed and accuracy with which analysts can identify and isolate network security threats.

References

[1] AfterGlow. http://afterglow.sourceforge.net/, accessed 9 Apr. 2010.

[2] C. G. Healey. Perception in visualization. Available at: http://www.csc.ncsu.edu/faculty/healey/PP/index.html. Last updated, 11 May 2009 (comprehensive review updated regularly).

[3] J. A. Jacko and A. Sears, editors. *The human-computer interaction handbook: fundamentals, evolving technologies and emerging applications*. L. Erlbaum Associates Inc., Hillsdale, NJ, USA, 2003.

[4] H. Koike and K. Ohno. SnortView: visualization system of Snort logs. In *VizSEC/DMSEC '04: Proceedings of the 2004 ACM workshop on Visualization and data mining for computer security*, pages 143–147, New York, NY, USA, 2004. ACM.

[5] G. A. Miller. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *The Psychological Review*, 63(2):81–97, Mar. 1956.

[6] Open Geospatial Consortium KML. http://www.opengeospatial.org/standards/kml/, accessed 21 Apr. 2010.

[7] Open Source Security Information Management (OSSIM). http://www.alienvault.com/community.php?section=Home, accessed 9 Apr. 2010.

[8] D. Shelly, M. Dunlop, R. Marchany, and P. Sforza. Using geographic information systems for enhanced security visualization. In *the 1st International Conference on Computing for Geospatial Research and Application (COM.Geo 2010)*, June 2010.

[9] D. J. Simons. Current approaches to change blindness. *Visual Cognition*, 7(1):1–15, Jan. 2000.

[10] Snort. http://www.snort.org/, accessed 2 Apr. 2010.

[11] I. D. Stroe, E. A. Rundensteiner, and M. O. Ward. Scalable visual hierarchy exploration. In *In Database and Expert Systems Applications*, pages 784–793. Springer-Verlag, 2000.

[12] S. M. Taylor. How much information is enough? decision-making and cognitive analysis. In *10th International Command and Control Research and Technology Symposium*, June 2005.

[13] RUMINT. http://www.rumint.org/, accessed 9 Apr. 2010.

[14] E. J. Wegman. Visual data mining. *Statistics in Medicine*, 22(9):1383–1397, Apr. 2003.

[15] J. M. Wolfe, K. R. Cave, and S. L. Franzel. Guided search: an alternative to the feature integration model for visual search. *Journal of Experimental Psychology: Human Perception and Performance*, 27(3):419–433, Aug. 1989.

[16] A. Woodruff, J. Landay, J. L, and M. Stonebraker. Constant information density in zoomable interfaces. In *Proceedings of the 4th International Working Conference on Advanced Visual Interfaces (AVI '98)*, pages 57–65, 1998.