Free visualization tools for security analysis and network monitoring

By Sam Abbott-McCune, A.J. Newtson, Robert Ross, Ralph Ware, and Gregory Conti

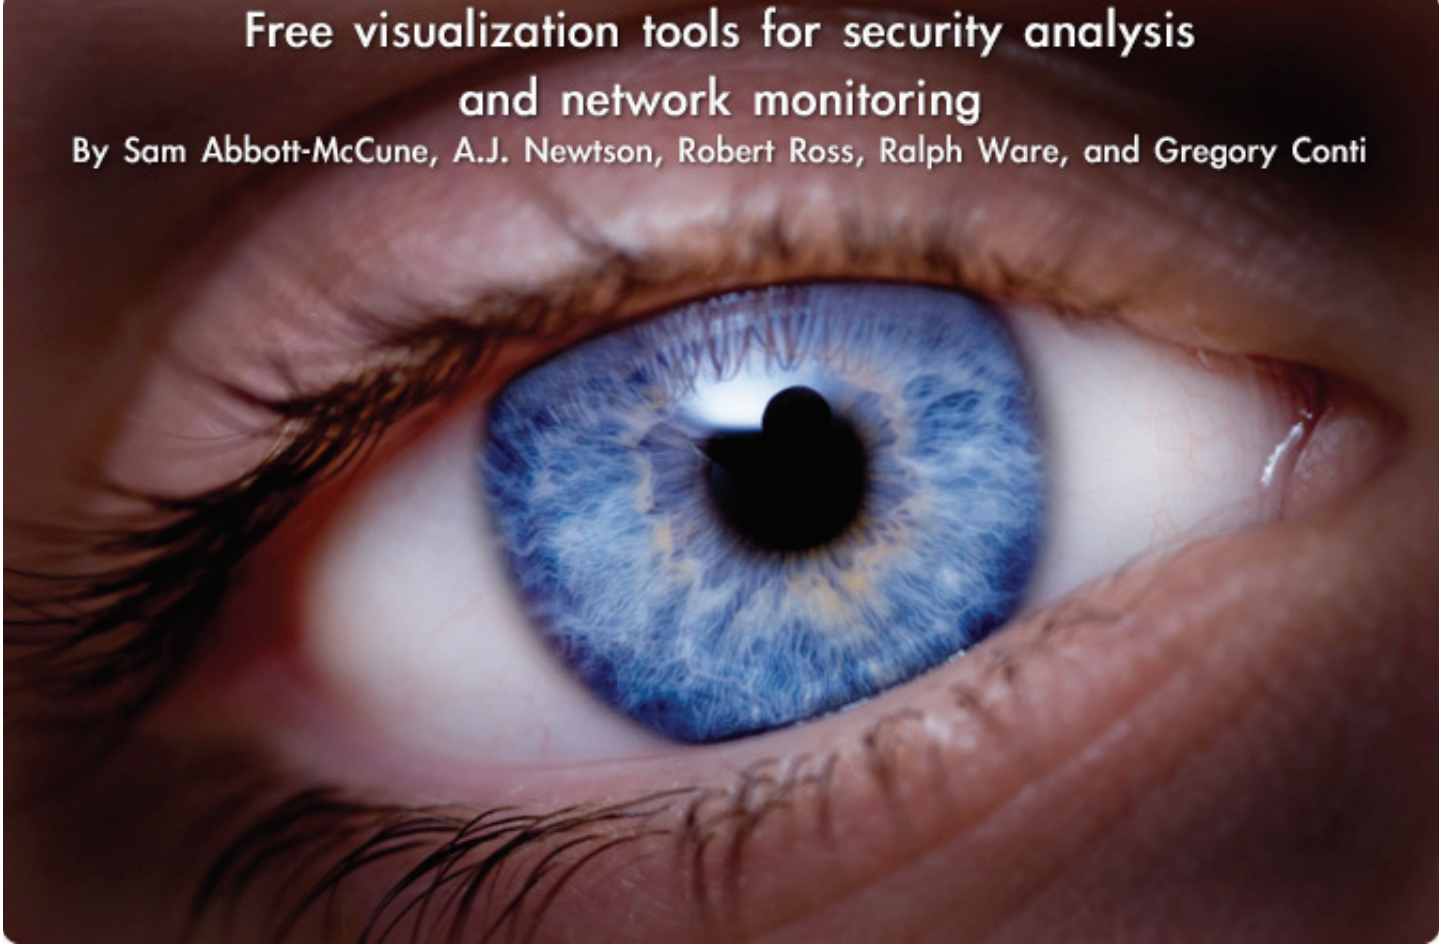

Whether you are a security analyst, system administrator or technical manager, chances are you are confronted with an overwhelming sea of security related data. Typically, we analyze this data with textual reports, command line scripts, or simple pie graphs and bar charts. However, there are much richer ways to analyze and explore the data using information visualization techniques. Information visualization systems attempt to create insightful and interactive graphical displays that exploit the human's extremely powerful visual system.

If done correctly, users will be able to examine more data, more quickly and see anomalies, patterns and outliers in ways that textual data simply cannot provide and machine processors cannot detect.

In this article, we present a number of free visualization systems that you can use to help find insight in your data. Where applicable, we've also included links to other tools you may wish to explore. In order to provide a broad overview of available options, we've sought out tools across a number of security related domains, including: network visualization, packet visualization, network management, and port scan visualization, as well as general purpose tools that can be used with many types of security data.

Network visualization

The Interactive Network Active-traffic Visualization (INAV), see Figure 1, is a monitoring tool that allows network administrators to monitor traffic on a local area network in real-time without overwhelming the administrator with extraneous data. The visualization tool can effectively perform a variety of tasks from passively mapping a LAN to identifying reoccurring trends over time.

Currently, INAV supports Ethernet, IP, TCP, UDP, and ICMP. INAV is implemented using a client-server architecture that allows multiple administrators to easily view network traffic from different vantage points across the network.

Once established, the INAV server passively sniffs data from the network and dynamically displays activity between different nodes on the network while keeping statistics on bandwidth usage.

The current state of the network is stored and broadcast to the different INAV clients. The INAV client uses an intuitive, lightweight graphical user interface that can easily change views and orient on specific clusters of nodes. Once a node on the network is selected, the client highlights any node that has sent traffic to or from that location. The client receives the current state of the network with a variable refresh rate that is adjustable to limit INAV generated communications on the network. Installation of the tool is straight forward and its operation is very intuitive. The INAV server runs on any Linux operating system with root privileges, while the client was developed in Java and can be run on most operating systems.

You can download INAV at inav.scaparra.com and a detailed white paper is available at inav.scaparra.com/docs/whitePapers/INAV.pdf. You may also wish to explore other network visualization systems including Afterglow (afterglow.sourceforge.net), Doomcube (www.kismetwireless.net/doomcube), Etherape (etherape.sourceforge.net), FlowTag (chrislee.dhs.org/pages/research/projects.html#flowtag), and Packet Hustler (shoki.sourceforge.net/hustler).

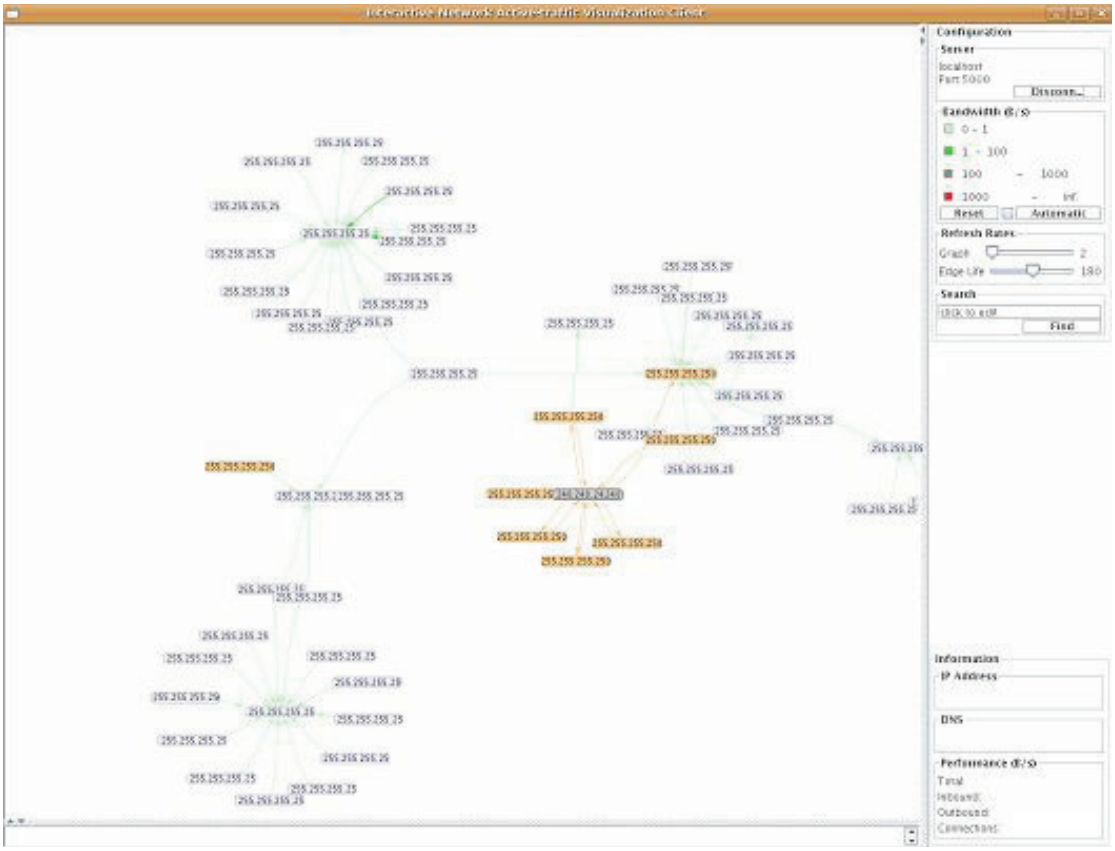

Figure 1: The Interactive Network Active-traffic Visualization (INAV) system passively sniffs network traffic and dynamically creates network graphs.

Nmap visualization

The fe3d network visualization tool, see Figure 2, is an open source application that works in conjunction with nmap and presents scan results using a 3-dimensional cone tree visualization (see citeseer.ist.psu.edu/308892.html for more information on cone trees).

Fe3d can be used with either imported nmap XML scan files or, alternatively, the user may launch and observe scans in real time. It also allows the user to routinely monitor network nodes for security issues such as open ports without requiring textual analysis. Fe3d gives the user the same scan results as command-line nmap, but in a very intuitive, easily understood 3-dimensional visual format by

graphically portraying the network node's operating system, IP address, and all open ports found on the node. This tool requires the following additional open source applications, Xerces-C++ XML parser (xerces.apache.org/xerces-c/install.html) and wxWidgets(www.wxwidgets.org/downloads/). We initially encountered difficulties interfacing the XML parser and wxWidgets on Linux operating systems, but found Windows installation to be quite straightforward, although we recommend that you use a recent version of Microsoft Visual C++ for easier installation. If interested in installing and testing fe3d go to projects.icapsid.net/fe3d. There you will also find very well written installation and configuration instructions.

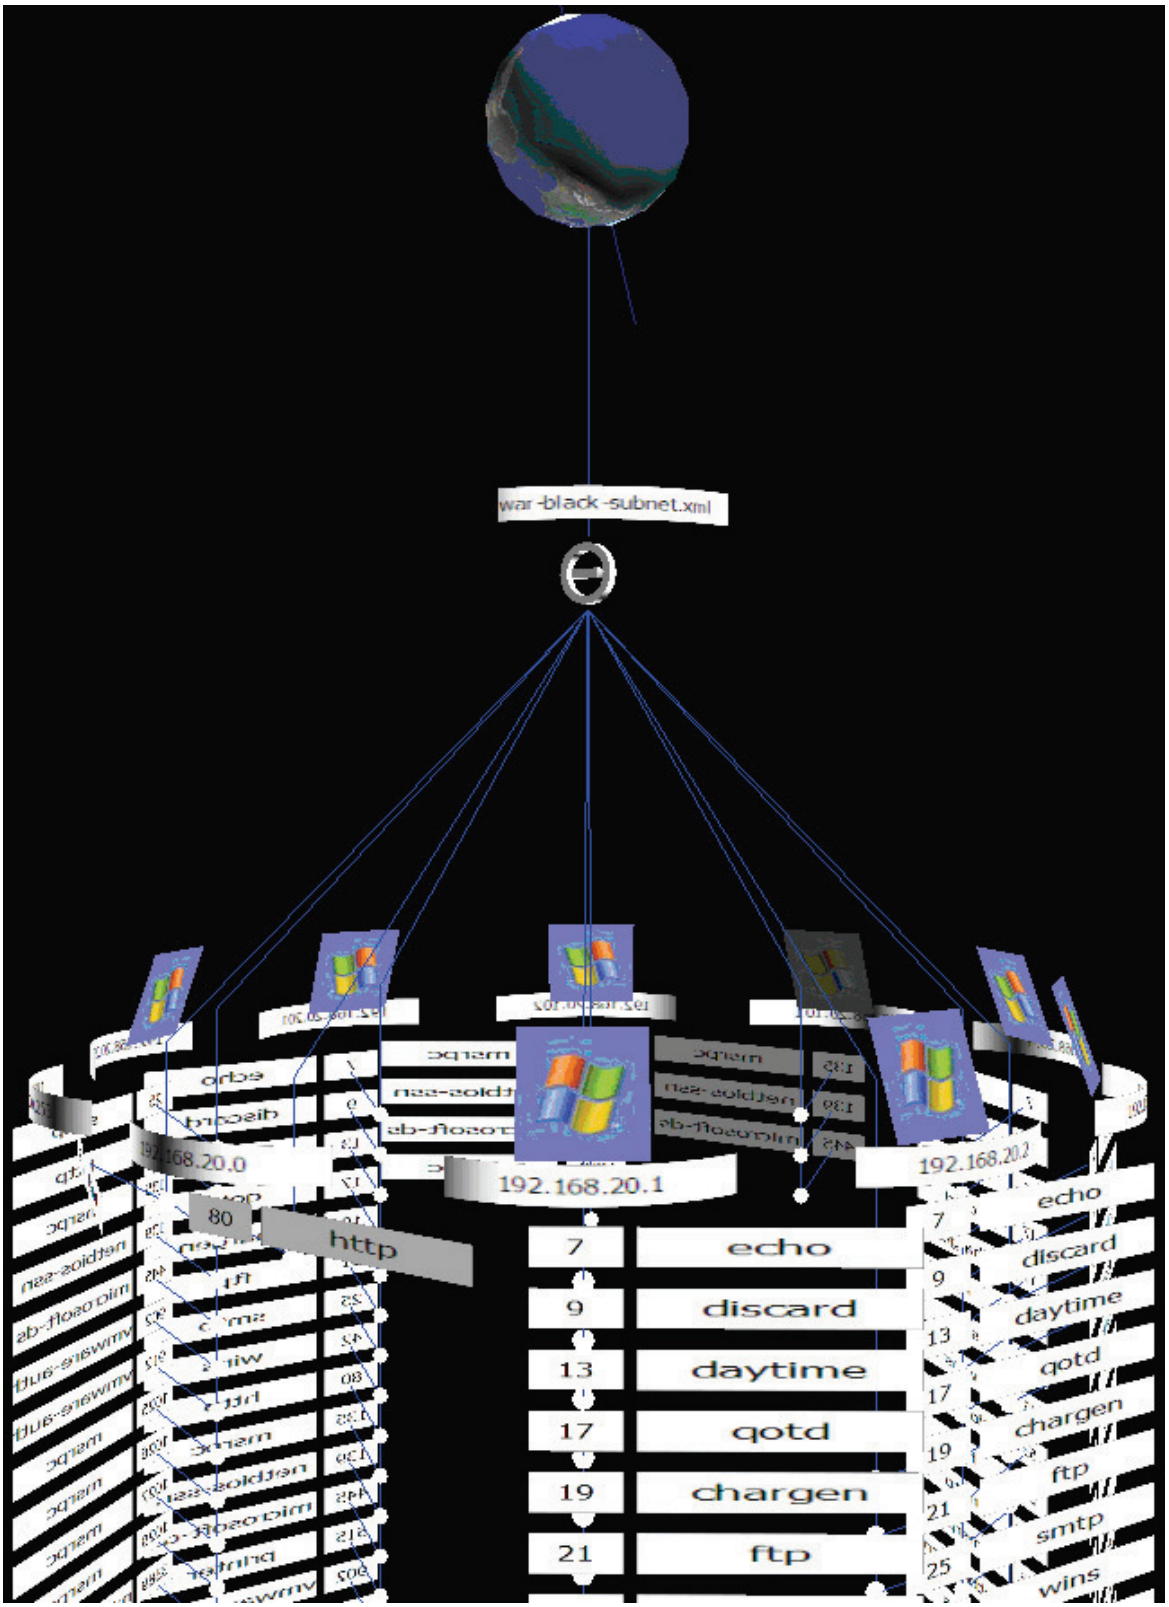

Figure 2: The fe3d visualization tool acts as a 3D front end for nmap scans.

Network monitoring

There are a wide range of tools for network monitoring that give a graphical overview of activity on the network. One of the original tools on the market was WhatsUp Gold (www.whatsupgold.com). WhatsUp Gold is a robust and scalable, but expensive monitoring system.

Although WhatsUp Gold is a quality product, we found that OPManager (www.opmanager.com), see Figure 3, provides most of the same functionality in addition to being available as freeware for network administrators of less than 10 critical systems. Available for Windows and Linux platforms, OPManager installs a password protected webserver on the designated host, which is accessible from any client on the network. Some of the OPManager's functionality includes: WAN monitoring, services monitoring (Web, FTP, SMTP, LDAP, DNS, and more), application monitoring (MySQL, Microsoft Exchange, among others), Windows Services monitoring (IIS, DHCP Server, Event Log), URL monitoring, server, and switch monitoring, among other functionality. The network status is clearly represented by numerous reports and customizable network displays. OpManager is fairly intuitive and easy to set up.

Another product to try is Nagios (www.nagios.org). Nagios is Linux-based and Firefox-friendly. However, Nagios can be difficult to setup initially, but if you are familiar with PHP include files (.inc), then subsequent networks can be easily configured. Nagios is also a web-based client/server package which gives near real time updates. Another software package that is worth checking out is OSSIM (www.ossim.net). OSSIM is a Linux-based solution which goes beyond simple monitoring by integrating software such as Snort and Nessus.

Figure 3: The free version of OpManager lets a network or system administrator monitor up to 10 hosts.

Packet visualization

Wireshark (www.wireshark.org) is the best of breed tool for protocol analysis and provides a powerful text-based GUI for analyzing network traffic captures.

RUMINT (www.rumint.org), a prototype graphical network sniffer, takes a different approach. It lets an analyst compare large numbers of packets, including both header fields and payloads, using seven different visualization windows.

Figure 4 shows a parallel coordinate plot (top left) that allows comparison of up to 19 packet header fields, a binary rainfall view (top right) which plots the raw bits from each packet and a text rainfall view (bottom left) which uses Unix strings-like functionality to display printable ASCII characters, one packet per horizontal row, as well as a detail view (bottom right) to see a single packet in hexadecimal and ASCII. Not shown are three additional visualizations, a scatter plot that plots any combination of packet header fields on a two-dimensional display, an animated visualization of packets emanating from ports and IP addresses, and a byte frequency visualization that displays a scrolling graph of bytes contained within each packet. RUMINT uses a VCR metaphor, where an analyst loads a packet capture file and "plays" back the packets in the visual displays. Because it is a prototype, RUMINT lacks the robust filtering and protocol parsers included with tools like Wireshark and is limited to 30,000 packets. It runs on Windows XP and later systems, but has been used successfully on Linux using Wine.

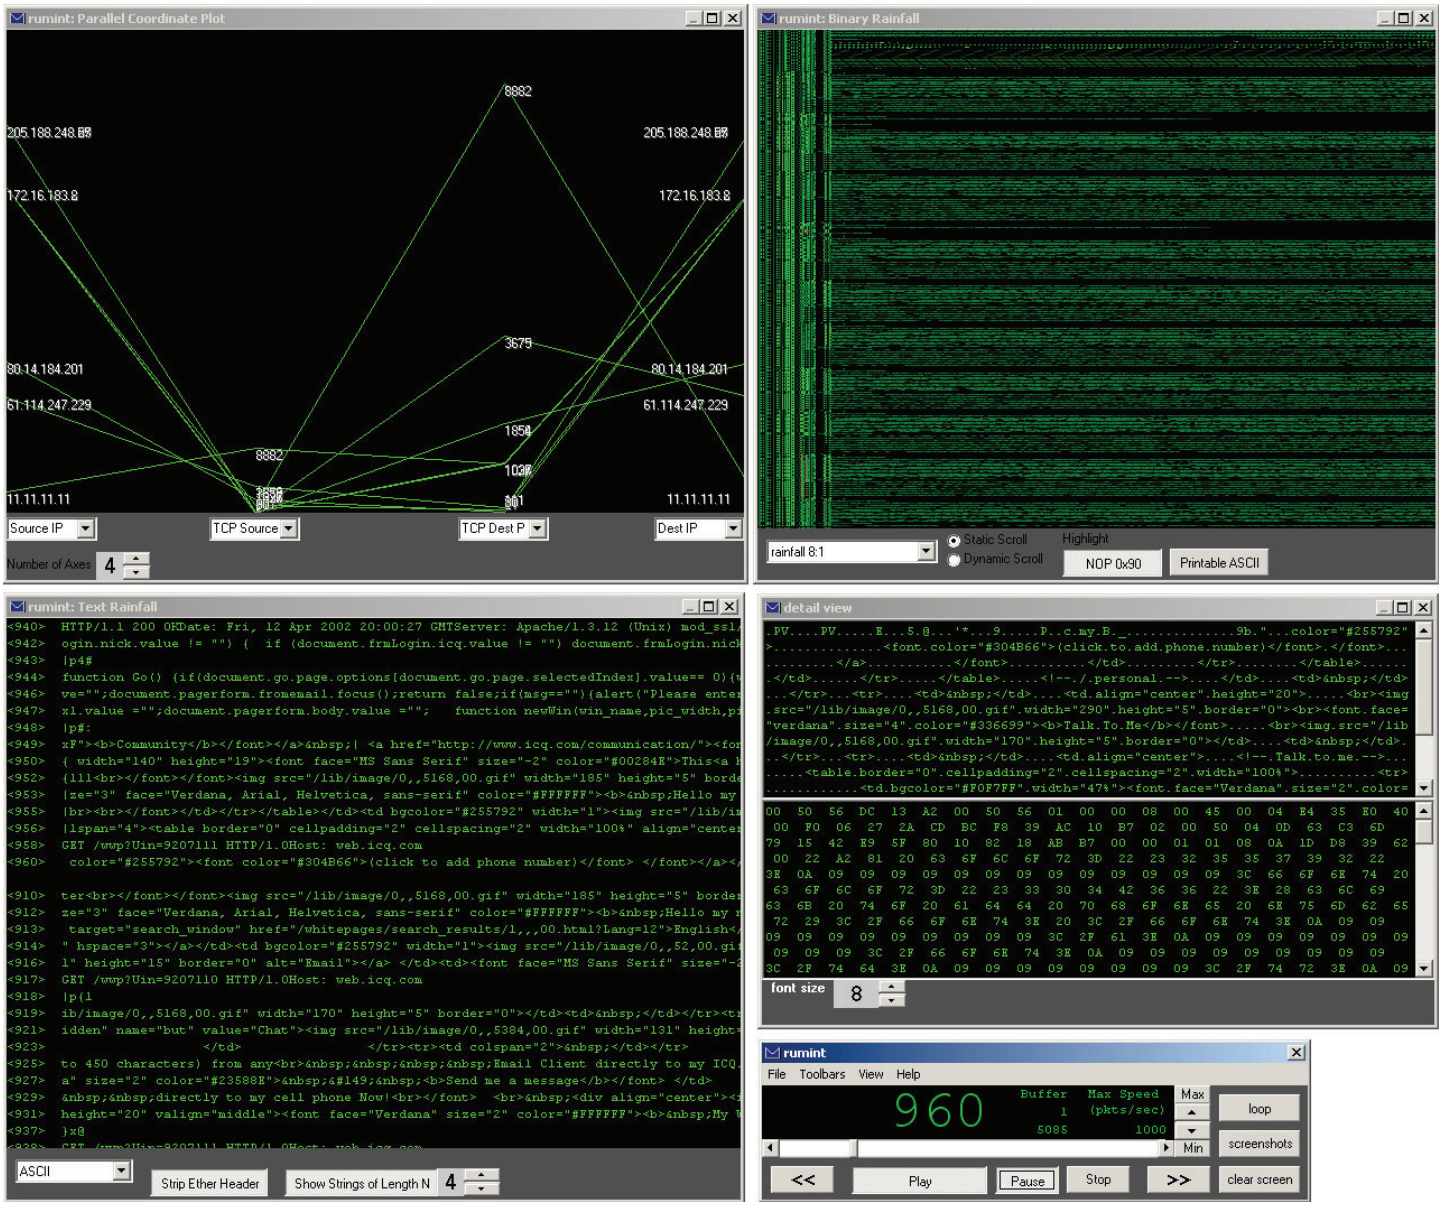

Figure 4: The RUMINT Visualization tool lets you capture and visualize network packets in real time.

General purpose visualization

Many Eyes is a free service offered by IBM and is an efficient and simple web-based application that incorporates numerous visualization techniques and facilitates collaborative analysis of security data. For example, after you collect network traffic from a tool such as Wireshark you can output the data to a comma separated value (CSV), upload it to

Many Eyes and view it using a number of interactive visualizations. (Note that a spreadsheet, such as Excel, can be very useful as an intermediate step to enhance or clean-up the dataset). A simple data table with named columns, each of the same length is required. Each column in the table supports two data types, text or numeric. You upload your data to Many Eyes via an HTML form by copying and pasting your data set.

Although Many Eyes has a dozen different types of visualization components, the network graph and treemap often provide the best insight into network traffic. Once a data set is uploaded to the Many Eyes server, you simply select a desired visualization component, allowing for flexible exploration.

Figure 5, is a snapshot of a network data capture from a Defcon Capture the Flag competition shown using the graph visualization component. The data set presented in this visualization contains the source and destination IP address of each packet. The Java applet is interactive and allows you to pan or zoom the view of the visualization as desired. Selecting a node, show in orange in the figure, highlights all adjacent nodes to facilitate analysis.

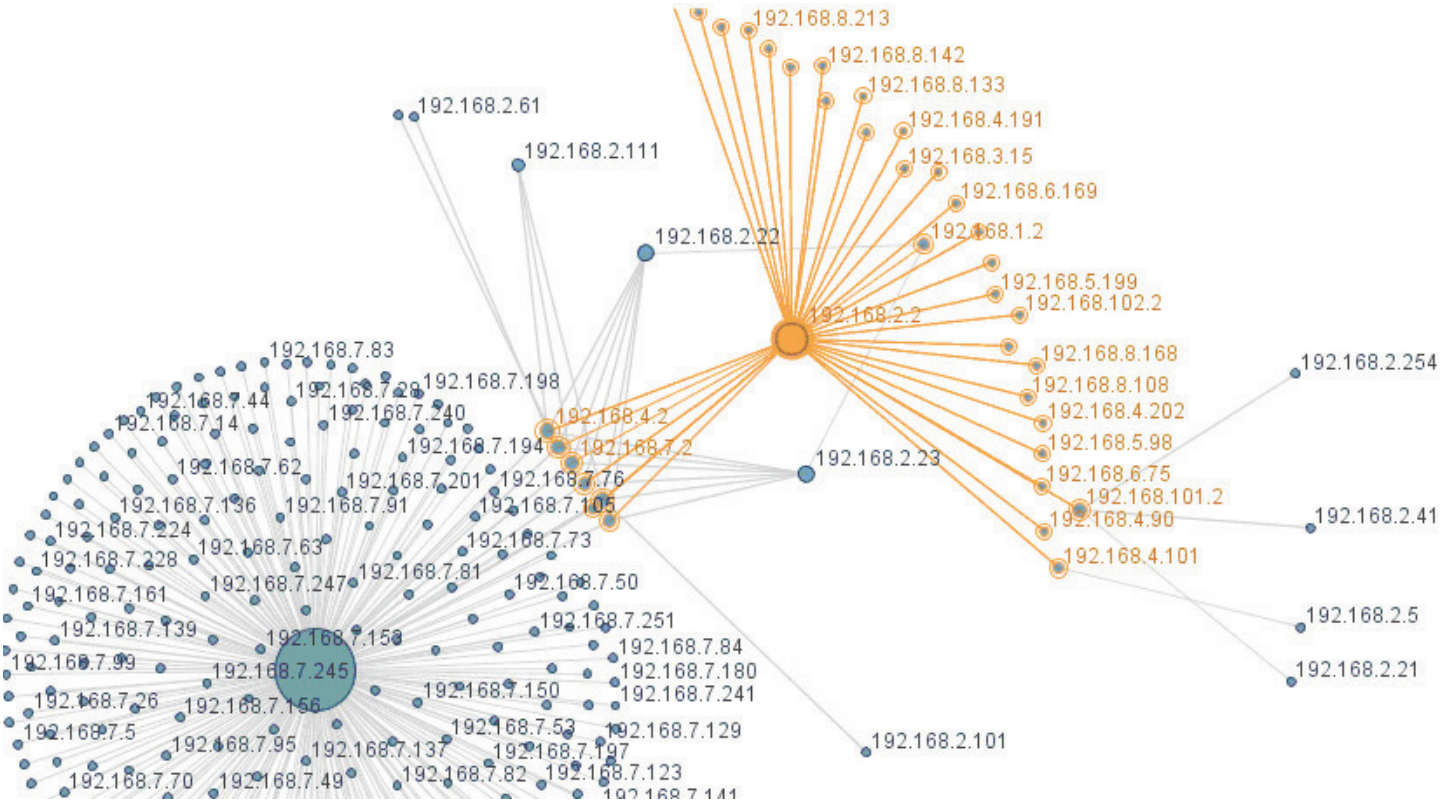

Figure 5: Using the Many Eye's visualization tool to graph a Defcon Capture the Flag Dataset.

On the following page is a snapshot of a similar network data capture, but using a treemap visualization technique. Treemaps are useful for visualizing hierarchical data, such as network addresses, as nested rectangles.

In the case of Figure 6 on the following page, the rectangles contain destination IP address, where the size of each rectangle corresponds to the quantity of packets, and the color corresponds to the destination port, where white is used for lower port numbers and dark orange for higher values. This visualization provides an alternative way to look at network data that can quickly identify patterns or anomalies, that a graph-based visualization cannot.

The benefit of Many Eyes is that it allows experimentation with a large number of visualization techniques and supports public collaborative analysis.

Registered users of Many Eyes (note that registration is free) can view, post comments and create additional visualizations based on a given dataset.

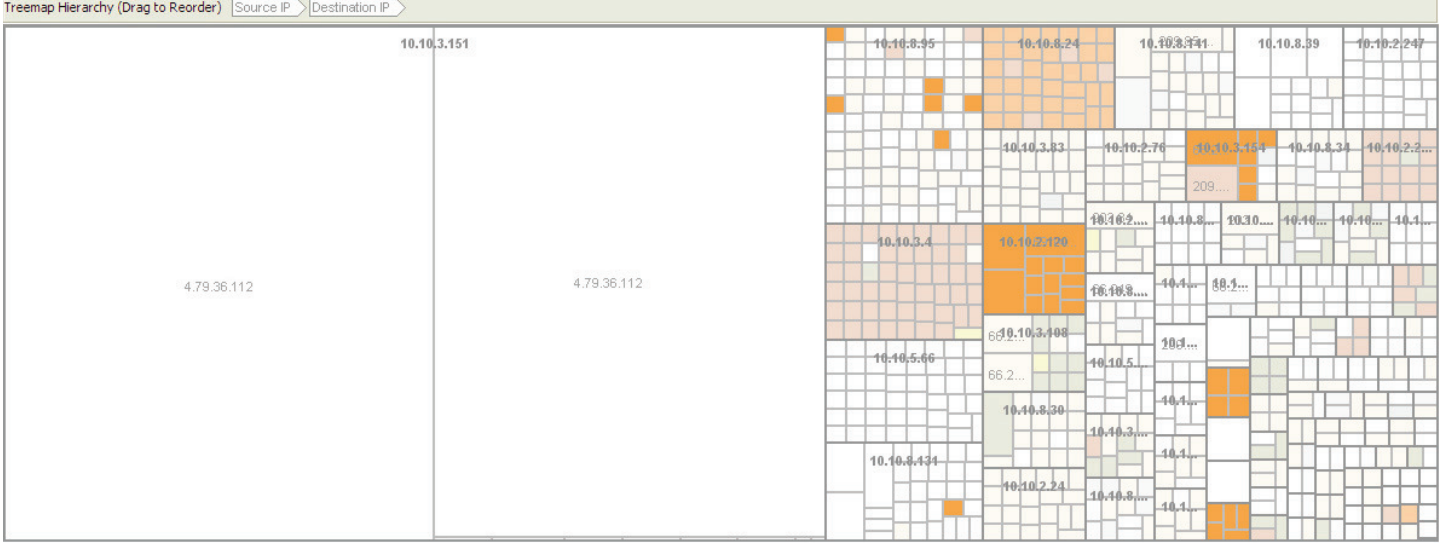

Unfortunately, at this time, there is no way to make a dataset or visualization private. Because of this issue, many network administrators may be reluctant to post data associated with their network. We leave it up to you to balance the risk of sharing your data against the strength of ManyEyes' visualization techniques and collaborative analysis facility. If you like Many Eyes, you may wish to explore other similar offerings such as Swivel (www.swivel.com) and the Google Chart API (code.google.com/apis/chart).

Name	Notes	URL
Afterglow	Graph visualization	afterglow.sourceforge.net
Doomcube	3D IP address and port visualization	www.kismetwireless.net/doomcube
Etherape	Network graph visualization	etherape.sourceforge.net
fe3d	nmap visualization	projects.icapsid.net/fe3d
INAV	Visualization of network bandwidth, source and destination nodes	inav.scaparra.com
FlowTag	Visualization of network flows	chrislee.dhs.org/pages/research/projects.html#flowtag
Google Chart API	Allows creation of dynamically generated charts	code.google.com/apis/chart
Many Eyes	General purpose visualization tool, accepts most CSV data	services.alphaworks.ibm.com/manyeyes
Nagios	Network monitoring	www.nagios.org
OpManager	Network monitoring	manageengine.adventnet.com
OSSIM	Network and security data monitoring	www.ossim.net
Packet Hustler	Network traffic visualization	shoki.sourceforge.net/hustler
RUMINT	Packet-level sniffing and visualization	www.rumint.org
Swivel	General purpose charting tool	www.swivel.com
Wireshark	Best of breed protocol analysis tool.	www.wireshark.org

Conclusion

Security data visualization is an active area of research. In the near future expect to see tools that not only present data in insightful ways, but also help bridge the gap between human analysts and machine processing.

Human time and attention are a precious resource

Researchers are currently developing tools that allow insights made by human analysts to be offloaded to machine processors.

A good example is a tool that facilitates analysis of a new malware variant and allows the analyst to immediately generate a Snort signature.

We encourage you to evaluate the tools listed here, see Table 1, but more are being developed frequently. Two places to monitor for the latest developments are www.secviz.org organized by Raffy Marty and www.vizsec.org sponsored by SecureDecisions (www.securedecisions.com). For the latest security visualization research consider participating in the annual VizSEC Workshop (vizsec.org/workshop2008). The next VizSEC will be held in Boston on September 15, 2008 in conjunction with the Recent Advances in Intrusion Detection (RAID) Symposium.

One final note, we are currently in the process of attempting to catalog all open source security visualization projects, current and historical, if you have a suggestion please feel free to send an email to gregory-conti@usma.edu. We will freely share the results of the survey with the security community.

Sam Abbott-McCune is currently an Instructor, teaching Information Technology, Network Systems Management and Theory and Practice of Military IT Systems, at the United States Military Academy at West Point. He received his Master's Degree in Computer Science from Virginia Commonwealth University.

A.J. Newtson is currently an Instructor, teaching Theory and Practice of Military Information Technology Systems, at the United States Military Academy at West Point. He received his Master's Degree in Information Technology Management from the Naval Postgraduate School.

Robert Ross is presently an Information Technology Instructor at the United States Military Academy at West Point. He received a Master's Degree in Computer Science from Monmouth University.

Ralph Ware is currently a Course Director and Instructor, teaching Information Technology, at the United States Military Academy at West Point. He received his Master's Degree in Computer Science from the Georgia Institute of Technology.

Gregory Conti, Director of the Information and Technology and Operations research center and Assistant Professor of Computer Science at the United States Military Academy, is the author of Security Data Visualization (No Starch Press) and the RUMINT visualization tool. His work can be found at www.gregconti.com.