

Visualization Techniques for Intrusion Behavior Identification

¹Dr. Robert F. Erbacher, *Member, IEEE*, ²Menashe Garber

Abstract – Current intrusion detection techniques are plagued with false positives and false negatives. Ensuring that intrusions are not missed requires that administrators filter through enormous numbers of false positives. In this work, we are attempting to improve the administrators ability to analyze the available data, make far more rapid assessments as to the nature of a given event or event stream, and identify anomalous activity not normally identified as such. To this end, we are exploring the roots of the identified activity, namely the underlying behavior of the users, hosts, and networks under the administrator’s auspices. We present here our work related to visualization as it applies to behavior and intrusion detection. We have found that the representations can be quite effective at conveying the needed information and resolving the relationships extremely rapidly.

Index terms – Intrusion detection, information visualization, anomaly detection, behavior analysis.

I. INTRODUCTION

One of the principal difficulties inherent to intrusion detection is the discrimination between innocuous and malicious events. This is made particularly difficult due to the high volume of false positives and false negatives that plague current analysis techniques. In attempting to improve the analysis process we must consider what makes an event malicious; namely the context of the event. It is the context of the event that ultimately identifies its overall purpose or behavior. Though often assumptions are made about particular types of events; e.g., unrequested packets. We can assume they are malicious without truly identifying their nature or intent. Generally, any event in and of itself is neither malicious nor innocuous. It is the event stream (or sequence of events) that identifies the overriding behavior and thus must be examined to derive whether an event is truly malicious or innocuous (beyond initial assumptions). The goal then must be to analyze the behavior of said event stream. These requirements for analysis of event streams is indicative of the success of low and slow probing, as the individual packets associated with such probes are not associated with an active probe or similar such event stream and thus are not flagged as rapid scans and their associated event streams are. This allows such attacks to progress where unsophisticated scans fail.

Such an analysis itself is difficult due to the obfuscation of the event stream. The knowledgeable hacker will attempt to distribute their attack both spatially and temporally, requiring extensive correlation of events to identify such obfuscated events.

The goal of this work is to provide visualization techniques to aid in the analysis of intrusion related data through the identification of behavior, notably changes in behavior (e.g., anomaly detection). This can be particularly useful for misuse as well.

In this paper, we will first examine the need for behavioral analysis in intrusion detection and how behavior is applied to such analysis. Second, we will examine how our techniques apply to the representation of behavior. Third, we will discuss the development of novel visualization techniques designed specifically for the representation and analysis of behavior. We then present examples of the application of the techniques showing their benefits. Finally, we will sum up and present tasks for future work.

II. BEHAVIOR ANALYSIS

The intrusion detection field currently suffers from enormous difficulty in effectively analyzing the available data. This difficulty is the result of the high number of false positives and false negatives generated. This in and of itself results from the lack of any consistent paradigm indicative of normal or abnormal behavior. It is very easy for an experienced hacker to form much of their activity as that of typical activity. This is further complicated by the general volume of traffic on networks these days. Additionally, many of the attacks that are identified are merely noise from scans being performed by script kiddies and the like. These types of scans are of far lesser concern due to the effectiveness of firewalls and traditional security measures at blocking and isolating them. So how can we identify the true (competent) attacks through the noise?

Our focus has been to revert back to the most fundamental questions related to the analysis of events:

- What makes an event malicious as opposed to innocuous?
- What differs between innocuous users and malicious ones?
- What is the goal of the malicious events?

¹ Utah State University, Department of Computer Science, UMC 4205, Logan, UT 84321, Robert.Erbacher@usu.edu

² University at Albany – SUNY, Department of Computer Science, LI67A, Albany, NY 12222, menashe79@yahoo.com

This concept of goal goes at the underlying foundation of the proposed techniques. Namely, we are attempting to identify the goal, the overriding context, or more specifically the behavior of the identified event stream. It is the behavior of the event stream that aids in identification of the true intent (goal) of the event and the intent of the initiator.

Additionally, behavior is an important indicator for misuse, particularly internal misuse. The idea with internal misuse is that we are identifying the behavior of an individual on the system as that of either an intruder or of a valid user not behaving in accordance with specified policies, especially security policies as they relate to computer and network usage. Identifying such perpetrators is critical for reducing the vulnerabilities inherent to the environment.

III. ANIMATED NETWORK REPRESENTATION AND BEHAVIOR

We have developed visualization techniques for the representation of network activity and the aid it can provide in identifying anomalous activity. Our goal when developing these techniques was solely that, to aid identification of anomalous activity. As we formalize our exploratory and analysis techniques, as they apply to intrusion detection, we are identifying the effectiveness with which such techniques apply to behavior analysis and their effectiveness at this task. These techniques are disadvantaged due to the fact that they only provide snapshots of activity at any one point in time. However, this can be critical as one piece of the puzzle in aiding analysts' differentiation between innocuous and malicious events streams. For example, the techniques discussed later in this article use this representation, as a detail view to garner explicit details as to the activity of a host or user identified as behaving unusually in the summary displays, providing a level of detail the summary displays cannot.

Figure 1 shows an example of the animated representation, showing multiple local and remote hosts. The clarity with which unusual behavior is exhibited can be extremely valuable. The image shows an example of a primary server as well as several workstations on the local network (bottom of the image). Many remote nodes (top of the image) are connecting to the local network, principally to the server. This is expected behavior. One connection, however, goes directly to a local workstation. The crossed lines and deviation in line angles aids this identification. This user is then identified connecting through to the server in a pattern indicative of intruder migration.

An additional visual artifact is exhibited by one of the local workstations (2nd from the left) as it is maintaining individual connections to a large number of local workstations simultaneously. Once such unusual activity is identified it can be explored to identify the reasoning behind the activity; perhaps through querying of the user. In any event, an appropriate response can be initiated.

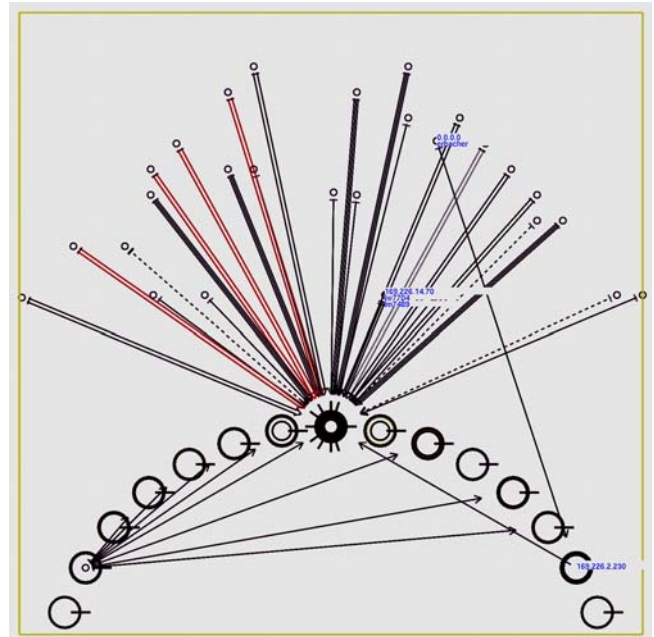


Fig. 1. Example at one point in time of a network monitoring visualization. Activity of the network (behavior) can clearly be identified and examined at that point in time. Probing is shown which allows examination of activity.

These visual artifacts rely on the notion of pre-attentive vision. Such displays provide overall behavioral representations of the network at any point in time through network monitoring of system log files and system statistics. The collation of log files in a distributed environment is critical for identifying such activity. Probing of the display provides details as to the activity, including hostname and usernames when available; to allow quick analysis and resolution of identified anomalies. This rapid analysis is critical in today's network environments in which administrators must examine and resolve enormous numbers of anomalies, all while attempting to perform traditional system administrative type duties.

In terms of scalability, the technique will similarly work effectively for larger numbers of hosts and servers, screen real-estate allowing. The key with larger networks is the fact that their will generally only be a single connection server, maintaining the philosophy as described above. Alternatively, many networks provide a more distributed connection paradigm, in which users connect to a wider range of workstations and servers from remote locations. This scenario isn't of as much concern since it intrinsically lacks safety and thus should not be deployed in critical environments and such deployments must assume a larger number of compromises. Additionally, the distribution of connections ensures that each system will have a fairly small number of connections, allowing said hosts to be monitored and examined efficiently.

A second example, showing the monitoring of a single host is provided in figure 2. In this example, the activity of multiple remote hosts is indicative of unusual behavior. The

bright red nodes are indicative of port scans identified by portsentry [12]. The rapid sequence of port scans can clearly be seen. Here a fade effect, akin to radar type displays, is critical. Even though these port scans are occurring from different remote hosts, the extensiveness and synchronization of the activity and thus the goal of the activity is clear, i.e., a single coherent attack divided and instigated from multiple machines simultaneously to inhibit detection.

These examples from our initial work show how clearly behavior can be visually represented. Such representation of behavior is ideal for the examination of intrusions and misuses due to the need to determine the goal and extensiveness of the underlying activity in order to determine the maliciousness and severity of the activity. The effectiveness of these techniques led to the realization that they are in fact aiding analysis of behavior. Thus, additional visualizations were developed to aid in the intrusion detection process with the specific goal of representing behavior.




Fig. 2. Visual representation of a single monitored host. Again the representation is designed to aid the examination and analysis of anomalous activity (behavior). The fade effect is critical for associating events occurring in rapid sequence.

IV. SUMMARY VISUALIZATIONS AND BEHAVIOR

The effectiveness of the animated visual representations led to the need for techniques more specifically geared towards the representation of behavior. In particular, one of the goals of these new techniques was to represent history information through static summary visualizations. Thus, we are attempting to visually represent entire time periods, and the encompassed event streams, such that the behavior indicative of the event stream is represented. The goal is to build off of the techniques described previously and

enhance these techniques through the application of additional visualization techniques that correlate events and identify the behavior more effectively.

A. Event→Threat Mapping

The following table shows a sampling of mappings from events found in typical log files and the resulting threat level. These mappings are adjustable by editing a simple text file. In this representation, one represents least threatening events while ten represents most threatening. Also of note is the ability to incorporate the results of other tools such as portsentry. No limits are placed on the tools that can be incorporated. All that is required is adjustment of the parsing engine to incorporate the new events.

It is important to keep in mind that we are examining the raw data in this scenario. The representation of raw data relies on the capabilities of the human user, both through their visual capacities, intuition, and prior experience to identify activity warranting further examination. We do include the results of other tools within the visualization environment (e.g., portsentry). However, the visualization environment itself does not perform any analysis of the data. The user is expected to perform further analysis. Ultimately, the results of any outside tool may be included, such as data mining techniques. Given the variety of current data mining techniques, their variance in applicability, and range of results [10] it is critical to use these results as additional inputs rather than solitary inputs such that the user can visually examine all data rapidly and make an assessment as to the threat of the activity. By correlating all available data, the user can make far better decisions and far more rapidly than typical tools and algorithms.

Syslog Identifier	Numerical Value
ALERT	9
ANONYMOUS	5
INETDFTP	2
INETDTELNET	2
LOGININCORRECT	8
PORTMAP	7
PORTSENTRY	10
PRIVILEGED	2
SUDO	10
TELNET	2

Fig. 3. Table summarizing threat levels for a subset of typical events.

B. Histogram Visualizations

A simple representation of behavior was developed and is presented in figure 4. This is a simple line-based histogram in which each remote host generates a single histogram line. Individual event streams can be selected and highlighted for more exact analysis. This technique's primary limitation is the occlusion of the activity exhibited at the bottom of the display. The occlusion is managed by allowing individual event streams to be selected from drop down lists, either for

the display entirely or related to an individual point selected by the user. In the latter case, all event streams passing through the selected point are listed for selection.

As an example of the need for behavior analysis, consider the line-based histogram shown in figure 4. In this example, we have a tight cluster of many high severity events, appearing on the right side of the display. When such activity is encountered the question as to whether this is an attack must be raised. If the system administrator had to rely solely on the textual data then this unusual behavior would likely have been completely missed, allowing a potential attack to go unchallenged.




Fig. 4. Line-based histogram representation of intrusion data. Time is shown on the X-axis and threat level (severity) is shown on the y-axis.

With the visualization environment, this unusual behavior sticks out like a sore thumb. Additionally, we have the ability to quickly analyze this activity to determine the meaning behind the user's behavior. In this scenario, what is actually occurring is that we have an academic environment in which an individual attempted to login to one of the public access computers multiple times. In fact, the individual attempted to log into one machine three times in a row (the standard time out parameter). The severity level indicated the failure of these attempts. The individual then attempted to log in to a second workstation, within the same lab, three times. A final attempt was made to login to a third workstation a single time. At this point the individual appears to have given up.

Clearly, knowing this is an academic environment; it is obvious we have a student that has forgotten their password. This can be removed from consideration as a serious threat. The key, however, is the rapidity and detail to which the user's behavior can be analyzed. This is a

critical ability for the intrusion detection community. Additionally, this example shows how the behavior and meaning of the behavior is critically dependent on the environment in which it occurs.

C. Pixel-Based Histogram Technique

When developing more effective techniques geared towards analyzing behavior, we desired techniques effective at immediately identifying typical activities of concern to system administrators, such as users sharing account information, intruded accounts, intruded systems, etc. Additionally, we desired techniques suffering less occlusion than the technique described in the previous section. Figure 5 shows an example of the developed capability. In this representation, activity of each remote host is represented as color plots, with each bar representing a single remote host. A green → red color scale is used to be representative of the threat level of an event. Additionally, the mode of the bar representation may be switched to have each bar represent a single local host or a single user.



Fig. 5. Summary display showing per host activity. Remote hosts are shown here. Variations from one host to another are observable, as well as variations within single hosts.

An additional histogram is overlaid on top of each pixel bar. This histogram is designed to complement the threat level presented in the pixel bar by representing activity in the form of connectivity information. Thus, the histogram can represent who is performing the connectivity (i.e., by username) when available, the host connected from, or the host connected to. Thus when taken together we can have a display in which each bar represents a single remote host. The color changes of the bar represent activity over time as related to the severity of the identified activity. The histogram overlay then can represent the host connected to, indicating whether the identified activity was all to the same system or to different systems. The complete set of

possible parameter mappings is represented in figure 6. The first column identifies the representation of each bar, does each pixel bar represent a remote host, local host, or identified user. The second column indicates reasonable modes for the histogram overlay to provide additional information to the user. Such information may be representative of the username, remote host, or the local system connected to.

Given the format of typical IP addresses, we apply a simple mapping metaphor to convert the IP address into a relative value for placement of the histogram values. This relies on the fact that we are more concerned with greater deviations than with smaller deviations. A user connecting in rapid sequence from the same remote subnet isn't likely an issue. A user connecting in rapid sequence from disparate remote subnets (e.g., from different countries) would be of critical concern. Thus, the mapping is based on the top-level network address value.

Bar Representational Modes	Histogram Representational Modes	
Local host	Remote host connected from	Username
Remote host	Local host connected to	Username
Username	Remote host connected from	Local host connected to

Fig. 6. Visual format representational modes. Given the primary mode for each bar the corresponding choices for the histogram is shown.

The goal is to have this histogram provide a representation of behavior such that either individually or in correlation the presented parameters will aid the analyst in identifying anomalous activity. The example in figure 5 shows the activity of a subset of the remote hosts, with the histogram representing local host connectivity information. This allows the analyst to identify connectivity patterns of remote hosts in conjunction with the severity level of the activity that has been initiated by that remote host.

Clearly, the resolution of the histogram prevents exact identification or differentiation of hosts or users. However, the goal is not exact identification of hosts but rather relative identification of hosts and identification of overall behavior, especially changes in behavior. Even more, the correlation between the threat level of the activity on a given system in conjunction with the connectivity activity can greatly aid analysis of the behavior of the networked environment; identifying activity needing further analysis; in essence making the resolution of false positives and false negatives of events falling under the auspices of the capabilities of these techniques far more rapid and efficient. This goes a long way towards aiding our goal of using behavior as a principal metaphor towards the identification and analysis of intrusions and misuses in a visual interface. In essence two representations of behavior are included. The first is the threat level of each event in a specified user

or host's event stream. The second is the systems involved in this activity, the duration of connections, and delay between connections. Figure 7 exhibits these characteristics. This figure shows several hosts' activities simultaneously in a zoomed view for clearer presentation. Merely examining the activity on each line gives the sense of consistency or deviation in behavior. For example, the fourth host maintains consistent activity throughout the recording period. With the lack of any threatening activity and the consistency of activity this how would not cause alarm?




Fig. 7. Zoomed example showing the activity of several remote hosts. As described in the text the first host shows consistently threatening activity. The second host shows isolated threatening activity. The fourth host shows consistently unthreatening activity. The last host shows highly unusual activity requiring immediate attention.

The second host shows similar activity throughout the course of the majority of the activity with a significance variation at one time point. This variation relates to the threat level of the event. Correlation of this event with connection activity can aid the administrator in rapidly determining the threat of the given activity as a whole. In particular the appearance of connectivity appearing immediately after threatening activity indicates this host would require further examination, particularly at this time point.

An additional sequence is shown in the first host. Here, normal activity is a continuous sequence of threatening events. Clearly, this isn't unusual for that host. It is likely that such activity would be examined at one time, determined to be acceptable and no further considerations given, and thus not wasting the administrator's time. Should the pattern deviate from this "normal" pattern that further analysis would be required.

The fifth host shows very unusual activity likely to be a focused and directed port scan of the system as large numbers of portsentry alerts are identified in rapid sequence. This is followed by more typical activity. This is a host that should receive priority attention to ensure a compromise did not succeed. Is the typical activity part of a break-in or the normal activity of a valid user? Such brute force port scans are generally well protected against. However, the extent of this port scan and the fact it did not occur on other systems raises concern. This correlation and differentiation is made simple with this correlated display.

These five analysis examples show the benefits of the developed techniques. Much of the analysis can be done directly within the visualization without the need for analyzing the text directly. This allows hosts to be classified as innocuous or questionable (and requiring further analysis

far more readily). This meets our goal of attempting to develop capabilities that improve the efficiency and effectiveness of analysts.

Currently the histogram goes to a zero level when no connectivity or activity is present and jumps to a representational level indicative of the active parameter when activity begins. An alternative representation would fade the histogram to the bar color, hiding it, when no activity is present. This alternative view may make changes in connection characteristics more clearly visible.

It should be noted that our data collection paradigm incorporates data collected with `inetd` run with the `-t` option, which reports initial connection requests before authentication. This can be valuable in identifying connectivity or attempted intrusions for which no username is associated. Given that our current paradigm is limited to host-based data this information is extremely valuable. We will likely identify far more benefits with network traffic data as well. This initial connectivity information has proven extremely valuable in identifying anomalous activity.




Fig. 8. Example of the summary display showing a per-user based representation. The activity of individual users is shown. Far less activity is visible due to the volume of activity occurring by unauthenticated network traffic.

This allows identification of numerous anomalous activities that otherwise would not be identifiable. Most attacks, especially those by competent hackers, of which we need to be particularly concerned, will be exhibited through changes in behavior and not one of the traditional approaches.

An alternative example is presented in figure 8 which shows a per user model of the technique. This example has far less detail due to the amount of activity occurring by

unauthenticated users. Of note in this example is the detail to which the activity of individual users can be identified, discerned, and comprehended. Comparing the user-based mode to the remote host-based mode clearly shows the amount of activity that differs and provides an analysis capability in and of itself. While there appear to be no truly severe attacks present, as are indicated in the remote host – based example, this display does show promise for analysis. In particular, we can easily identify the activity of individual users and examine said activity for potential violations, attempts at migration, and insider threats.

These techniques provide a novel and effective improvement for the visual representation and analysis of IDS related data. Given the limitations of the current techniques, the effectiveness of these techniques and the progress being made should prove valuable both to practitioners and researchers in the field.

D. Summary Display




Fig. 9. Summary display showing overview of all activity represented in pixel-based histogram representation. The currently displayed region is bordered in blue. The selected element is marked with a red tag.

Given the limitations on screen real-estate, it is clearly impossible to show all activity of all hosts simultaneously. Consequently, we provide an overview display which provides a summarization of the activity of all hosts. This display uses a substantially reduced display representation that allows all elements to be represented but forgoes the inclusion of details as to the activity of each host, figure 9. Without these details it becomes difficult to differentiate or analyze the hosts individually. However, this representation does aid navigation of the hosts that are to be presented in the fully detailed display; e.g., the pixel-based histogram. In essence, the display represents a host as a one pixel high bar. No overlay is provided but threat level is still

incorporated. This provides for context and focus while interacting with the environment.

This display is interactive and linked to the fully detailed display. The hosts presented in the fully detailed display are bordered in blue. The currently selected host is identified through a red tag. The selection rectangle may be dragged or moved to select any range of hosts. This allows for rapid examination of all hosts within the data set. The fully detailed display may be panned directly as well, this summary display is updated as it does so.

V. IMPLEMENTATION

The discussed techniques are provided within a single environment implemented in C++ and using OpenGL. Tcl/Tk is used for the user interface capabilities. By providing a single environment we essentially provide coordinated views with the disparate views displayable simultaneously. Interactions and selections are carried across views to improve analysis efficiency.

VI. RELATIONS TO PREVIOUS WORK

This work is impacted by the work ongoing in multiple fields. This includes other behavioral research, both from the psychological and the human computer interaction fields. Intrusion detection research, which is credited with providing the foundations of what we are trying to represent, identify, and comprehend. As well as visualization research, which greatly aids in creating more effective visual representations of the data in such a way as to allow more effective examination by the user?

A. Behavioral Security Research

Behavior has been considered with respect to security from a variety of viewpoints. At the fundamental level, understanding how users behave aids us in better understanding how we can implement more effective security protocols. The Human Computer Interaction (HCI) field has been examining the needs of users in order to develop user interfaces that better match the expectations and typical responses by users. For example, Adams et al. [1] analyzed the effectiveness of typical password policies. Stanton et al. [15] have created a taxonomy of behaviors and the relationship of these behaviors to security policies.

More recently, others have begun exploring the applicability of behavior to more varied aspects of security. Behavior analysis, also called anomaly detection, attempts to identify malicious activities through the identification of changes in behavior and goes to the foundation of many recent data mining intrusion detection tools [4], [8]. Such tools, however, suffer from high false positive and false negative rates [10]. This work begins to explore the visualization of behavior for intrusion detection.

B. Intrusion Detection Research

In addition to the behavior-based approach discussed previously, other heavily focused areas of intrusion

detection research has focused on the application of data mining [9], neural networks [11], and signature-based (snort [14]) techniques for the identification of intrusions. Many additional techniques have been explored but much less extensively; a full survey is beyond the scope of this article. While, these techniques do suffer severe limitations they are noteworthy due to the progress they have made in attempting to identify intrusions. Even with this progress these techniques suffer limitations. Data mining and neural network techniques suffer from lack of accuracy due to the chaotic and noisy nature of the data source, thus the high rates of false positives and false negatives. The signature-based techniques are only applicable to known techniques and can easily be circumvented by a competent hacker.

C. Visualization Research

Visualization has been applied successfully to the analysis of many data sources. Its application to intrusion detection, however, has to this point been limited. Work that has been applied has proven successful in aiding the examination and comprehension of the available data. However, many challenges remain, specifically in dealing with the scale of the data, the number of parameters, the chaotic nature of the data, and the need to provide clearly distinguishable artifacts for all appropriate anomalies, attacks, intrusions, misuses, etc.

Example techniques include the intrusion detection techniques by Teoh et al. [16], Scott et al [13], Wood [17], as well as the work by Erbacher et al. [6]. The work by Teoh et al. examines the effectiveness of visualization for the analysis of Internet routing data and the applicability of such data to intrusion detection. Scott et al. explored simple node and link visualization techniques with the application of haptics. Wood describes basic graph-based visualization techniques, such as pie charts and bar graphs, and how these techniques can be applied to typical network data available to all system administrators. This work provides a fundamental description of how visualization can be implemented and its application to such data, as well as the meaning behind the identified results. The work by Erbacher et al. provides extensive application of visualization directly to typical intrusion related data for the visualization and correlation of data.

Additionally, many techniques have focused on network monitoring. These network-monitoring techniques have direct applicability to intrusion and attack detection but so far have not been extended to provide the level of capability needed. Such techniques include: immersive network monitoring [7], E-Mail usage analysis [5], bandwidth utilization [2], and web access statistics [3].

Currently available visualization techniques are thus not suitably designed for the intrusion detection task and challenges. Those tools that are designed for intrusion detection are generally designed to support very specific subtasks; as with the routing data analysis by Teoh et al. mentioned previously. Our work is designed to expand the

visualization capabilities to provide more robust general capabilities and tools to aid administrators and analysts in the analysis of IDS related data.

VII. CONCLUSIONS

We have examined techniques through which behavior can be visualized for the purpose of identifying intrusions, misuses, and attacks within a networked environment. The visualization techniques are appropriate for identification of traditional attack signatures as well as changes in behavior. These techniques provide a visual solution to the identification and analysis of changes in behavior as opposed to a purely algorithmic one. The environment was also designed in to support selection between multiple parameters both as principal and secondary dimensions. This allows selection between displays incorporating remote hosts, local hosts, users, and threat levels.

Ultimately, the goal is to aid analysts in identifying the true goals and behavior of events and event streams much more rapidly and effectively by relying on the capabilities of the human visual system and the human analysis process. This should greatly reduce the number of false positives and false negatives consuming valuable time and resources. Additionally, we presented examples of how typical attack metaphors will appear within the environment.

VIII. FUTURE WORK

The current techniques represent the threat level of individual events, forcing the analyst to analyze the impact on the event stream as a whole. We must examine the true impact of examining individual events versus event streams. This will become particularly important when we incorporate examination of network traffic data. With network traffic data each event in and of itself represents a much smaller unit. This greatly complicated the analysis process, as we must provide the capabilities for analyzing, filtering, and comprehending far more data.

IX. REFERENCES

- [1] Adams, A., Sasse, M. A. & Lunt, P. (1997) "Making passwords secure and usable" in H. Thimbleby, B. O'Connell & P. Thomas (eds.), "People & Computers XII (Proceedings of HCI'97)" Springer, pp. 1-19.
- [2] Richard Becker, Stephen Eick, and Allan Wilks, "Visualizing Network Data," Readings in Information Visualization: Using Vision To Think, Stuard Card, Jock D. Mackinlay, and Ben Shneiderman, editors, Morgan Kaufman Publishers, pp. 215-227, 1999.
- [3] Tim Bray, "Measuring the Web," Readings in Information Visualization: Using Vision To Think, Stuard Card, Jock D. Mackinlay, and Ben Shneiderman, editors, Morgan Kaufman Publishers, pp. 469-492, 1999.
- [4] Dorothy E. Denning, "An Intrusion-Detection Model," IEEE Transactions on Software Engineering, Vol. SE-13, No. 2, February 1987, pp. 222-232.
- [5] Stephen G. Eick and Graham J. Wills, "Navigating Large Networks with Hierarchies," In Visualization '93 Conference Proceedings, San Jose, California, pp. 204-210, October 1993.
- [6] Robert F. Erbacher, Kenneth L. Walker, and Deborah A. Frincke, "Intrusion and Misuse Detection in Large-Scale Systems," Computer Graphics and Applications, Vol. 22, No. 1, January/February 2002, pp. 38-48.
- [7] Mike Fisk, Steven A. Smith, Paul M. Weber, Satyam Kothapally, and Thomas P. Caudell, "Immersive Network Monitoring," In Proceedings of the Passive and Active Measurement Workshop, 2003.
- [8] Trent Henry, "Securing the Enterprise with Network Behavior Anomaly Detection," Research Report, The Burton Group, October 2003.
- [9] W. Lee, Salvatore J. Stolfo, and Kui W. Mok, "A data mining framework for building intrusion detection models," IEEE Symposium on Security and Privacy, pp. 120-132, 1999.
- [10] John McHugh, "Intrusion and Intrusion Detection," International Journal of Information Security, Volume 1 Issue 1 (2001), pp 14-35, 2001.
- [11] S. Mukkamala, A. H. Sung, "Learning machines for Intrusion Detection: Support Vector Machines and Neural Networks," Proceedings of International Conference on Security and Management, pp. 525-531, 2002.
- [12] David Sarmanian, "Deploying PortSentry - A Simple and Free Barrier From Inside Hackers," SANS Institute, GIAC GCIA Practical, January 2001.
- [13] Craig Scott, Kofi Nyarko, Tanya Capers, and Jumoke Ladeji-Osias, "Network intrusion visualization with NIVA, an intrusion detection visual and haptic analyzer," Information Visualization, Vol. 2, No. 2, pp. 82-94, 2003.
- [14] Roderick W. Smith, "Network Monitoring with Snort," Linux Magazine, May 2003.
- [15] Jeffrey M. Stanton, Cavinda Caldera, Ashley Isaac, Kathryn R. Stam, Slawomir J. Marcinkowski, "Behavioral Information Security: Defining the Criterion Space," The Systems Assurance Institute, Syracuse University, Syracuse, New York, 2003, <http://sai.syr.edu/facultypapers/Stanton%20-%20BehavioralDomain.pdf>
- [16] Soon Tee Teoh, Kwan-Liu Ma, S. Felix Wu, "A Visual Exploration Process for the Analysis of Internet Routing Data," Proceedings of the IEEE Visualization Conference, IEEE Press, 2003.
- [17] Alex Wood, "Intrusion Detection: Visualizing Attacks in IDS Data," SANS Institute, GIAC GCIA Practical, February 2003.