Methods of Visualizing Temporal Patterns in and Mission Impact of Computer Security Breaches[1]

Anita D'Amico, Ph.D. and Mark Larkin
Secure Decisions, a Division of Applied Visions, Inc.
AnitaD@avi.com

Abstract

The primary objectives of this project were to design and prototype 3-D visual representations of: 1) time trends in information security events; and 2) mission impact of information security events. Secure Decisions, a Division of Applied Visions, Inc., interviewed several information security analysts in the US Dept. of Defense and in commercial industry and compiled a list of information requirements to meet these objectives. We then designed several temporal and mission impact visual displays, and prototyped several of them in 3-D Virtual Reality Modeling Language (VRML).

Future development of mission impact displays will require that mission impact data be collected and stored in a form that can be accessed by visual displays. Secure Decisions found no readily available techniques for automatically determining mission impact, or even for determining which network resources are required by specific mission-critical tasks (i.e. mission x device dependencies). To help fill this gap, Secure Decisions outlined the types of data that an automatic system would have to generate to provide users with an efficacious display of mission impact of security incidents.

1. Introduction

Situational awareness refers to an operator's ability to see the "big picture", i.e. an overall understanding of a system's status and any impediments to the system's ability to achieve its goals. Studies on situational awareness [1] [2] [3] [4], in general, conclude that the operator who is situationally aware must have:

- *perception* (aka recognition) of the elements of the current system state;
- *comprehension* (aka assessment) of how all the individual elements of the system state fit together and affect the system's ability to achieve its goals; and
- *projection* (aka prediction) of the future system state based on either a continuation of the present state or a change in state brought about by purposeful actions.

Cyber defense situational awareness requires the information assurance analyst to:

- Accurately perceive the overall security state of the information infrastructure;
- Comprehend the relevance of current and past security incidents and their impact on the organization's mission; and
- Project the effects of both unmitigated security incidents and the courses of action (COA) that may be taken to counteract those incidents.

To achieve situational awareness the information assurance analyst must form a mental model of the information security state and use this model to assess new information and project future status. Analysts often find that knowledge of previous security incidents helps them to assess the nature and sophistication of the current threat, the timing of the attack, and the next likely steps in the attack sequence.

However, aids to achieving situational awareness are limited in the relatively new field of information assurance. While previous research in other fields has shown that visual representations can be useful aids to decision-making under time pressure [5], the information security field has used visual representations sparingly. Currently

[1] The work described in this paper was supported by a Phase 1 Small Business Innovative Research (SBIR) contract awarded by DARPA – Contract # DAAH01-00-C-R056

0-7695-1212-7/01 $10.00 © 2001 IEEE

available security tools are good at providing data, but they stop short at providing an integrated picture to the user. For visual displays to facilitate situational awareness they should provide the information security analyst with the ability to integrate data from many sources, correlate that data and see the big picture of the infrastructure's security state.

The Defense Advanced Research Projects Agency (DARPA) as well as other government agencies have repeatedly identified a need for visual representations that can help information security analysts to form a mental model of the past and current security incidents and to project the impact of those incidents on the ability to achieve mission goals. With this in mind, DARPA requested that Secure Decisions, a Division of Applied Visions, Inc., design visual representations that would assist the information assurance analyst in achieving cyber defense situational awareness. The work was performed under a Phase 1 Small Business Innovative Research (SBIR) project, entitled "Visual Representation of Cyber Defense Situational Awareness."

The two primary objectives of the project were to:

1. Design and prototype visual representations that depict time trends in information security incidents and changes in security state over time; and
2. Design and prototype visual representations that depict the impact of security breaches on mission-critical tasks. Within the context of this second objective Secure Decisions also defined the types of data needed to support an assessment of mission impact. These types of data will have to be collected and stored by future systems that support cyber defense situational awareness.

In our designs of visual representations that would support cyber defense situational awareness, we tried to use 3-D visualization to complement, not substitute, for the value derived from text and tabular displays.

2. Collection of Information Requirements

To achieve the primary objectives, Secure Decisions first studied how today's military and commercial information security analysts use information about the timing of incidents, the tools they use to gather and study temporal patterns, and how they assess mission impact of security events. Among those we interviewed were representatives from: the National Security Agency (NSA); Joint Task Force for Computer Network Defense (JTF-CND); the Global Network Operations and Security Center (GNOSC) at the Defense Information Systems Agency (DISA); the Department of Defense Computer Emergency Response Team (DODCERT); commercial financial institutions; and a commercial company that monitors the security of other commercial companies. Approximately ten very experienced information security analysts were interviewed in an informal, face-to-face setting. They were asked what types of aids they currently use to discover time patterns in and mission impact of cyber security breaches. They were also shown samples of visual aids and asked to comment on their potential utility. From these discussions and our own knowledge of the field, Secure Decisions developed specific requirements for what should be included in visual representations of temporal trends and the impact of security breaches on mission-critical tasks. We also defined the data domain that would be needed to fulfill these requirements and that had to be represented in visual scenes.

It should be noted that in our discussions and compilations of requirements, we used the term "security event" to mean any vulnerability, suspicious activity or actual breach that constituted a real or potential threat to the user's high-value information resources.

2.1 Temporal Display Requirements

With respect to detection of temporal patterns, we found that analysts currently use standard charting tools, like Microsoft Excel spreadsheets, to keep counts of the number of various security breaches over time. Some users are interested in trends in time measured in hours, while others were looking at trends over months. In general, we concluded that temporal visualizations should have the following capabilities:

- User-selectable time gradations (e.g. seconds, minutes, hours, days, months)
- User-selectable time range (e.g. from May 1 through June 15)
- User ability to annotate time grid (e.g. "June 13 – Checkpoint firewall vulnerability becomes public.")
- Ability to relate specific security events to time (e.g. show specific times that various probes occurred)
- Ability to relate the characteristics of security events to time (e.g. show the times at which certain types of attacks are most prevalent)
- Ability to relate target characteristics to time (e.g. show time periods during which specific operating systems or locations were attacked)
- Ability to relate attack sources to time (e.g. show period of time when certain attacker IP addresses are active)
- Ability to simultaneously relate types of security incidents, targeted resources and attack sources to specific time periods (e.g. depict the exact time and the order that specific workstations were probed, show both the operating system and location of the targeted workstations, and highlight any known information about the attack source)
- Depict frequencies of specific classes of incidents
- View sequence of incidents irrespective of absolute time (e.g. at Hanscom site #125, these events occurred in sequence from May 1-7)
- Depict duration of events (length of Denial of Service attacks on February 6-12)
- Simultaneously compare patterns of events over multiple user-specified time ranges (e.g. compare number of probes during April 1-7, May 1-7, June 1-7)
- Show time lapse between exposure (e.g. insertion of a vulnerability) and a related exploit
- Show differences between two user-selected times (e.g. show differences in vulnerabilities on a specific network on April 1 and June 1)

2.2 Mission Impact Requirements

Because no analyst we interviewed was regularly assessing mission impact of security breaches, we could not use their current methods as a baseline for gathering information requirements. Instead, we developed a few sample mission impact scenes early in the project, and used those scenes to generate discussions about the types of information users needed to analyze and how to best present it. We asked analysts what they considered their missions to be; how they would describe their missions, tasks within missions, and network resources they rely on; and what they thought would be useful to see on a mission impact display.

Interestingly, the use of the term "mission impact" display was confusing to many and elicited interpretations as varied as "displays that show the relationship between a power failure in Seattle and the U.S.'s ability to wage nuclear war" and "displays that show how the loss of Microsoft Word affects our ability to issue reports to field agents." As a result, we found that we gained more clarity and common understanding by substituting the term "functional impact" or "impact on mission-critical tasks" for the term "mission impact."

We also found that we had to distinguish between various types of network resources, such as hardware devices, software applications, databases, network services and connectivity. We categorized information resources into three major categories:

- A "Device" is a hardware platform used for information technology. It could be a workstation, printer, router, etc.

- A "Simple Resource" is a single application, database, service or file provided by a single device. A simple resource resides on one device. However, a single device could provide just one resource (e.g. it hosts personnel files for the entire organization) or a single device could provide many resources (e.g. it hosts word processing applications, accounting applications, and budget data for a specific department)
- A "Compound Resource" is more complex and represents a service to the organization (e.g. e-mail service, web access). A compound resource requires one or more devices and simple resources, and even other compound resources, to provide its service.

After interviewing potential users and clarifying our terminology, we concluded that mission impact visualizations should have the following capabilities:

- Illustrate all dependencies between network resources and mission-critical tasks
- Highlight dependence of a specific mission-critical task on network resources – i.e. show all the specific network resources that are required for a specific mission-critical task
- Highlight resource to missions dependencies – i.e. show all the mission-critical tasks that depend on a single specific network resource
- Provide user with ability to select the level of granularity he or she wishes to see regarding the dependencies between mission-critical tasks and network resources – i.e. collapse and expand across network devices, simple resources, compound resources and mission-critical tasks
- Show strength of dependencies (low, medium, high) between resources and mission-critical tasks
- Show "and/or" dependencies between resources and mission-critical tasks (e.g. to generate an Air Tasking Order one needs the Joint mapping application, and either access to the imagery database and a printer, or access to a secure fax machine)
- Show redundancies and substitutability of resources needed to support mission-critical tasks
- Depict how the strength of a mission-critical task's dependence on specific network resources varies based on the phase of a mission (e.g. the mapping application is only needed in the first phase of planning, whereas access to situation reports is needed throughout the entire planning process)
- Depict the sequential order in which specific resources are needed for mission-critical tasks (e.g. imagery files must be accessed by users before mission planning packages are put together).

3. Design and Prototyping of 3-D Displays

We designed several candidate visual scenes that captured as many of these requirements as possible and from those candidates we selected a subset for 3-D prototyping.

In designing our displays, we used 3-D space and numerous visual attributes of geometric objects to carry meaning in the visualization. Among the visual attributes we used were: shape, position, motion, size, dynamic size changes that express growth or shrinking, orientation, color, transparency, texture and blinking. We conformed to principles of good design, such as using gray as the background, so that the information conveyed by visual attributes of the objects in the foreground would stand out. [6]

One of the most compelling attributes of a 3-D visualization is the concept of dynamic perspective. Unlike 2-D representations, which have a single (front) perspective, 3-D visualization has infinite perspectives. In other words, since the objects in the scene are 3-D representations themselves, they can be viewed from the front, back, left, right, top, and bottom as well as any position in 3-D space. We used dynamic perspective and the navigation controls used with it in

our design of the temporal and mission impact displays.

The culmination of the project produced a demonstration of VRML (Virtual Reality Modeling Language) prototypes that could be manipulated in three dimensions. The visualizations were developed in VRML, and can be viewed using a VRML viewer such as Intervista WorldView (which can be licensed through Computer Associates). Although all of the visualizations shown in this paper were rendered and viewed using WorldView, other products, such as the Cortona VRML client from Parallel Graphics, can also be used to render and view 3-D displays.

An application program was developed in C++ on a Pentium platform to convert temporal and mission impact data in a test database into the VRML visualizations.

4. Sample Results

4.1 Temporal Scenes

To capture the need for detecting time patterns in security incidents, Secure Decisions designed four temporal visual display formats and prototyped three of these in 3-D VRML.

Figures 1 through 4 depict various views of the temporal event wall scene. The temporal event wall scene can be configured to meet all of the requirements identified by potential users except for the need to show differences in security events between two user-selected time periods. This is an important requirement that can be captured in different scene formats.

In Figure 1 the vertical wall displays information about: 1) time, for which the user specifies the range (e.g. January 1-10, 2000) and granularity (e.g. days, hours, seconds); and 2) event type (e.g. classes of vulnerability, attack, probes). The horizontal "floor" in the foreground is a "host grid" that provides information about the characteristics of the targeted hosts (e.g. their operating system, location, organizational affiliation). The association lines show the specific times that each event occurred on specific hosts. While this example depicts events as occurring in discrete points in time, it is possible to use the time wall to show duration of events, and the relative sequence to each other.

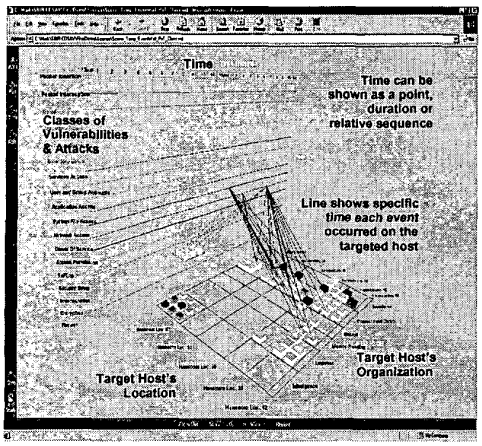

Figure 1. Front view of temporal display showing time that specific security events occurred, and the targets of their attacks

Figure 2 shows the rear plane or rear "floor" of the display. The rear plane provides information about the characteristics of the attack sources (e.g. IP address, number of hops it used to reach the target, etc.). It can also be used to show information about which specific sensors detected events, and the times that those sensors detected the activities.

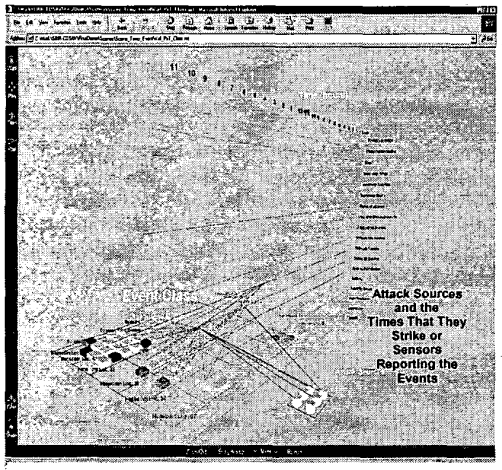

Figure 2. Rear plane of temporal display showing times that specific attackers are active

Figure 3 shows a top-down view. The user can simultaneously view the attack sources, the types of activities they engaged in, and the targets of their attacks – all related in time.

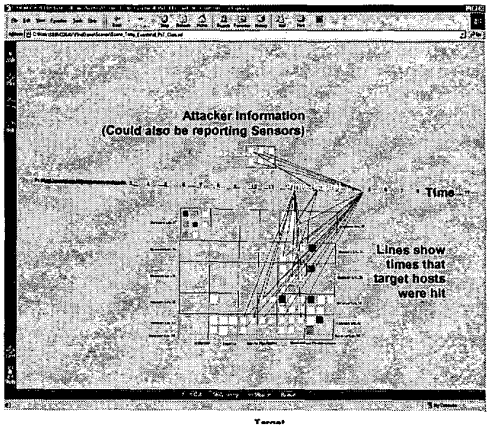

Figure 3. Top down view of temporal display showing how attackers, targeted hosts and events are related

Figure 4 depicts a different version of the frontal view. In this version, frequency distributions are shown on the vertical wall. The wall is divided into columns of time slots; in Figure 4 the time slots refer to minutes. The wall is also divided into rows of categories of security events (e.g. denial of service activities; access permission problems).

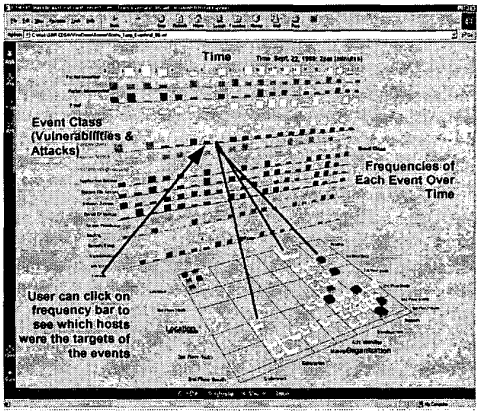

Figure 4. Frequencies of security events displayed by the time of detection and intended target

As a specific type of event is recorded in each time slot, the frequency bar increases in height. The user can click on the frequency bar to get more information. Upon clicking on the bar, association lines connect the frequency bar of events to the hosts that experienced the events. The user can then see what categories of hosts were attacked at specific periods of time, and the specific types of events they experienced at those time periods.

We do not anticipate any significant problems with capturing data to populate the temporal event wall. Most information security sensors (e.g. vulnerability scanners, firewalls, intrusion detection systems) store and report time-stamps with the alerts that it produces. Secure Decisions already has methods for interfacing visual displays with repositories of time-stamped sensor data, and those methods can be transferred to populating the temporal event wall.

4.2 Mission Impact Scenes

To capture the need to understand the impact of security incidents on the success of specific mission-critical tasks, Secure Decisions designed five mission impact formats and prototyped three of them in 3-D VRML.

It should be noted that the population of mission impact scenes is not nearly as straightforward as the temporal scenes. For mission impact scenes to be implemented in a real world environment, it will be necessary for the users to collect and store information about the interdependencies between network devices, cyber resources and the mission-critical tasks and missions they support. As part of another effort, Secure Decisions is working with DARPA's Mission Model Working Group to identify the specific types of information to be collected and the database schema needed to store that information.

Figure 5 depicts a sample of a mission dependency ring scene. The mission dependency ring format provides a basic structure that meets most of the identified requirements. It falls short primarily in the area of time phasing. The dependency ring format does not meet the requirement to show how the strength of a dependency varies with phase of a mission, not does, and it does not provide a means for depicting the sequence of cyber resources needed to support a mission-critical task.

In Figure 5, five rings are depicted:

Network device ring – This ring or disc is comprised of individual objects, each of which represents a specific device on the network. They can be workstations, routers, printers, etc. Devices are represented as cubes and occupy a single layer. Each device is labeled with its name. Drill down capabilities will allow a user to click on a network device and obtain additional information about its IP address, administrator, etc.

- Simple Resources – This layer is comprised of the simple resources provided by each device. A single device could provide just one resource (e.g. it hosts personnel files for the entire organization) or a single device could provide many resources (e.g. it hosts word processing applications, accounting applications, and budget data for a specific department). Three different object shapes are used to indicate the type of resource represented: a cube denotes an application program; a cylinder denotes data; and a sphere denotes peripheral devices.
- Compound Resources – This layer is comprised of resources that are more complex and represent a service to the organization (e.g. e-mail service, web access). These resources combine one or more devices and simple resources, and even other compound resources, to provide their service. Compound resources are arranged in one or more rings above the simple resources. Compound resources are represented as either a diamond shape or a cone shape. A diamond shape indicates the resource is an AND type. An AND type compound resource requires all its resource dependencies. A cone shape indicates the resource is an OR type. An OR type compound resource requires only one of it dependencies. Compound

resources can have other compound resources or simple resources as dependencies. When this happens, the compound resource is elevated to a higher ring.

- Mission critical tasks – This layer is comprised of objects each of which represent specific tasks that must be achieved by the organization (e.g. ATO generation, production of mission situation reports, shipping of supplies)
- Missions – This layer represents major goals that the organization has to achieve. Each goal requires multiple tasks to be accomplished for the goal to be reached.

The lines that connect objects in each ring represent dependencies. Although Figure 5 does not show variations in line thickness, the scene can be manipulated to show stronger dependencies using thicker lines and weaker dependencies using thinner lines.

In Figure 5, the user has requested to see all the dependencies between cyber resources and the mission-critical tasks and missions they support.

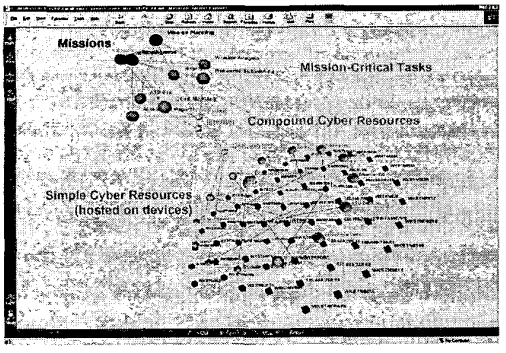

Figure 5. Mission impact scene showing dependencies between missions, the mission-critical tasks that support the missions, and the cyber resources needed for the mission-critical tasks

Figure 6 combines the concept of a mission dependency ring with a host grid, similar to the host grid seen in the temporal event wall. The bottom or "floor" of the display consists of a grid of devices that are on the network. The top of the display shows resources, tasks and missions that require those devices. In this example, the user has clicked on a specific device on the network that may have sustained an attack. The user wishes to see the specific cyber resources (e.g. applications, databases, services) that are hosted on that device and the specific mission-critical tasks that need those resources for their successful completion.

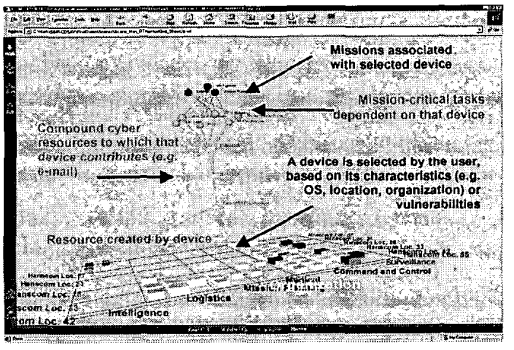

Figure 6. Mission impact scene showing what cyber resources and mission-critical tasks will be affected by a breach of a specific device

Secure Decisions identified several critical information requirements for building visual aids that will enhance the information security officer's situational awareness. We specifically focused on visual aids to improve the detection of time patterns in security breaches, and to improve the assessment of mission impact of security breaches. In our next phase of work for DARPA, Secure Decisions will refine these concepts and develop a fieldable prototype of a subset of them. Those prototypes can then be used to test their value in improving the situational awareness of information security analysts.

6. References

[1] http://www.reticular.com/Library/SA/sys_arch.html

[2] http://www.jas.co.jp/crm/english/decision.htm

[3] Endsley, M., "Toward a theory of situation awareness in dynamic systems", Human *Factors*, *37*, 1995, pp. 32-64.

[4] Fracker, M., "A theory of situation assessment: Implications for measuring situation awareness", In *Proceedings of the Human Factors Society 32nd Annual Meeting*, Human Factors Society, Santa Monica, CA, 1988.

[5] Coury, B. and Boulette, M., "Time stress and the processing of visual displays", *Human Factors, 34*, 1992, pp. 707-725.

[6] Tufte, Edward R., *Envisioning Information*, Graphics Press, Cheshire Connecticut, 1990.