

# Visualization as an Aid for Assessing the Mission Impact of Information Security Breaches<sup>1</sup>

Anita D'Amico, Ph.D. and Stephen Salas  
Secure Decisions, a Division of Applied Visions, Inc.  
[AnitaD@SecureDecisions.com](mailto:AnitaD@SecureDecisions.com)

## Abstract

*One of the objectives of the DARPA Phase 2 SBIR project entitled "Visual Representation of Cyber Defense Situational Awareness" was to prototype 3-D visual representations of mission impact of information security events. Secure Decisions, a Division of Applied Visions Inc., prototyped several mission impact visual displays, and incorporated the most promising ones into a visualization software architecture called SecureScope™ that can be deployed in an operational setting. The mission impact visualization system can interface to any common relational database that links IA events to the cyber assets (e.g. workstations, data files) on which those events occur, and includes relationships about the organization's mission-critical tasks that depend on those assets. It visually represents "bottom-up" mission impact analyses showing the mission effect of a specific breach on a specific asset, as well as "top down" analyses showing which cyber assets must be secured for supporting the achievement of certain mission-critical tasks.*

## 1. Introduction

Not all cyber security breaches are equal. Some are more important than others, because some breaches affect the ultimate mission of an organization more than others. Because it's impossible to protect everything, organizations need to know where to concentrate their Information Assurance (IA) resources and which IA events they should respond to first. People responsible for assuring the availability and integrity of mission-critical assets need to know which cyber assets are most critical to their organization and need to rapidly assess whether security events on those assets are likely to significantly impair the mission of their organization. Helping people to see how breaches of various assets affect their organization's mission was the primary focus of the mission impact visualizations developed by Secure Decisions under DARPA-sponsored Phase 1 and 2 SBIR efforts (Contracts DAAH01-00-C-R056 and DAAH01-01-C-R044).

The identification of assets that are deemed mission-critical and breaches that are deemed most harmful varies with each organization. Therefore, the application

of IA resources should vary between organizations. In many business organizations, IA resources are concentrated on protecting cyber assets and responding to security events on those assets that are perceived as significantly affecting revenue generation. By contrast, in the military world, IA resources are focused on those cyber assets responsible for the timely achievement of military objectives such as deploying troops, conducting surveillance, and launching attacks. Because of these differences in what assets are important to protect based on each organization's mission, one expects an e-commerce business to concentrate its information security devices and personnel on averting attacks on the company's commerce-critical servers, while a military mission planning organization may apply its resources to protecting the data links used to transmit Air Tasking Orders (ATOs).

While it may seem obvious to expect an organization to assign IA resources and respond to security events based on their mission impact, our experience indicates that organizations do not have a good understanding of which of their cyber assets are mission-critical, and how attacks on any of those assets will affect one or more of their organization's primary functions.

## 1.1 State-of-the-Art in Mission Impact Analyses

Our research found that there are few procedures and even fewer tools for conducting analyses that will yield information about the impact of a security breach on an organization's mission. Two main approaches to mission impact analyses are typically taken: the "bottom up" and the "top down."

**1.1.1 Bottom-Up Analysis.** In the bottom up analysis, the analysts' goal is to determine whether the unreliability or unavailability of a specific cyber asset is likely to deleteriously affect any of the missions of the organization. The analyst starts with a specific asset or capability (e.g. a database, data file, application, hardware platform, communication link, e-mail capability), and ultimately tries to determine which organizational missions would be affected if that asset or capability became unreliable or unavailable. In the course of doing this analysis, the analyst considers the

<sup>1</sup> The work described in this paper was supported by a Phase 2 Small Business Innovative Research (SBIR) contract awarded by DARPA – Contract # DAAH01-01-C-R044. The title of the project was "Visual Representation of Cyber Defense Situational Awareness – Phase 2".

types of events causing the unreliability or unavailability of the cyber asset; such as whether it is the result of a malicious act such as purposeful file corruption or denial of service, a security officer's intentional removal of the asset from service in order to avert a malicious act, or a non-malicious maintenance problem. He or she also considers dependencies among the cyber assets, i.e. whether other assets or capabilities incorporate or rely on the functionality of the questionable asset. For example, a single application program such as Outlook must be operational for some users to effectively use the e-mail capability; therefore the corruption of Outlook affects both the availability of the Outlook application and e-mail capability to some users. The analyst must also have an understanding of what mission-critical tasks and functions must be performed in order for an organization's missions to be achieved. For example, for an airborne attack mission to be successful, many tasks (e.g. select targets, plan mission, ensure aircraft has been maintained, load weapons) and sub-tasks (e.g. analyze potential target imagery, predict weather conditions, calculate fuel requirements) must be appropriately performed. Armed with all this information, the analyst can then hypothesize that if a particular server is corrupted, certain other cyber capabilities of the organization are affected, which in turn degrade the staff's ability to perform several specific tasks, which are needed by specific missions.

**1.1.2 Top-Down Analysis.** In the top-down analysis, the analysts' goal is to understand which cyber assets must be secured or assured for a specific mission to be successful. The mission may be quite global in nature (e.g. find all terrorist cells and eliminate them) or more specific to a particular department, agency or sector of the organization (e.g. complete intelligence report by November 30). The analyst starts with a specific mission and decomposes it into its tasks and sub-tasks, and then determines what aspects of the information infrastructure must be secured or assured to support those tasks and sub-tasks. Thus, the analyst might find that there are specific assets, such as workstations, databases, data files, and application programs, as well as specific capabilities such as reliable internet connectivity, that are crucial to achieving the mission.

**1.1.3 Mapping of Assets to Missions.** To conduct mission impact analyses, the analyst also has to have a basic understanding of which assets in the information infrastructure are needed to support specific mission-critical tasks in the organization. For example, an analyst needs to know that in order to perform the critical function of preparing an intelligence analysis, the person preparing the report must have access to various cyber assets and capabilities, (such as uncorrupted databases

containing imagery, text files containing message traffic, application programs for analyzing imagery, parsing messages, and formatting reports), and secured communications capabilities for transmitting and receiving intelligence reports.

Unfortunately, we found that very few organizations have a readily-available list of which assets in their information infrastructure are needed to support mission-critical functions; and among those organizations that do have such lists, they are rarely kept in a relational database that would provide analysts with easy access to its content.

## 2. Designs and Examples of Mission Impact Visualizations

Secure Decisions designed and prototyped several candidate visual scenes to meet the requirements for visual aids for mission impact analyses. We then incorporated the best designs into the SecureScope visualization system. To demonstrate the mission impact visual scenes, we interfaced the SecureScope system with a database of mission impact data that we drew from the Grand Challenge Problem scenario.

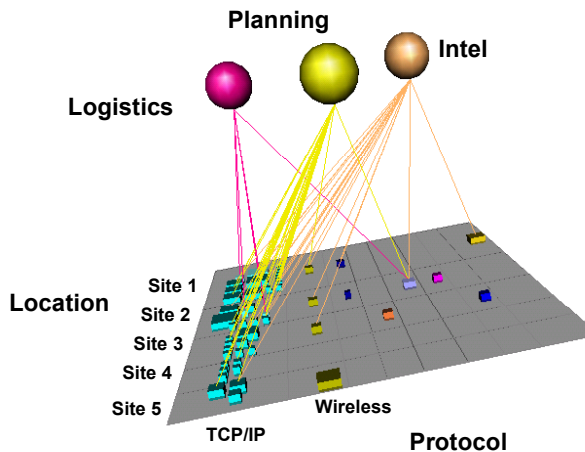
All of our designs utilize 3-D space through which the user can navigate. One of the most compelling attributes of a 3-D visualization is the concept of dynamic perspective. Unlike 2-D representations, which have a single (front) perspective, 3-D visualization has infinite perspectives. In other words, 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 that support this concept in our design of the mission impact displays. While all of our scenes use 3-D visualization, we strive to use 3-D to complement, not substitute for, the values derived from text and tabular displays.

All of the mission impact designs use geometric objects to represent entities in the mission impact database. For example, a cube might represent an alert from an intrusion detection system (IDS), or a cyber asset such as a workstation, while a sphere might represent a mission-critical task. By drawing lines from the cube to the sphere we depict a relationship or dependency between an alert and a mission-critical task, or a cyber asset and a mission-critical task.

The geometric objects are positioned in 3-D space in a meaningful way. For example, cubes that represent cyber assets might be positioned on a grid based on their physical location and the organization to which they belong. Other objects that represent alerts might be positioned on a wall based on the time they were detected. The geometric objects are assigned visual attributes to carry additional meaning in the

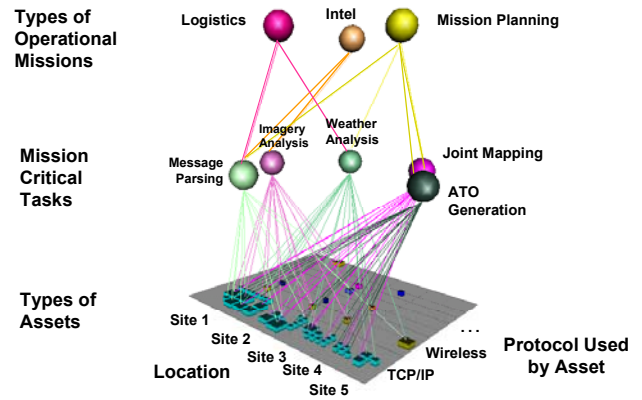
visualization. A cube used to represent an alert from an IDS could be colored red, orange or yellow to represent the severity of the alert, and could be made semi-transparent to represent the age of the alert. In other scenes, a pyramid shape could be used to represent a cluster of workstations that are all required for a particular mission. Blinking or motion could be added to the pyramid to indicate that one or more of the workstations in that critical cluster have experienced a security breach. Our selections of geometric objects, visual attributes, and 3-D positioning were chosen to conform to principles of good graphical 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. [2]

In the example in Figure 1, the boxes on the grid represent cyber assets or cyber capabilities such as workstations, application software, e-mail connectivity, etc. The sizes of the boxes indicate their level of criticality, and their color indicates the protocol they use. The boxes are distributed on the grid based on their location and their protocol. In this particular example, protocol is represented twice: by color and by position within a row on the grid. The spheres in the air indicate each of the mission-critical functions that must be performed. Lines drawn between the boxes on the grid the spheres in the air show associations between assets and mission functions. Those cyber assets required by the Logistics operations are connected via an association line.



**Figure 1 - Visual display showing relationship of cyber assets and the mission-critical operations they support**

The analyst can use a simple association scene, such as that shown in Figure 1, to perform either top-down or bottom-up analysis. Starting with the Logistics sphere, the analyst can see the specific assets that are needed for Logistics operations. In a similar fashion, the analyst can select a specific box on the grid that represents a critical asset and highlight those mission operations that depend on it.




**Figure 2 – Mission-critical operations at the top of the scene are related to the types of tasks and the cyber assets on which they depend.**

In Figure 2, another layer has been added between the cyber assets and the mission-critical operations. This new layer of spheres represents the specific types of mission-critical tasks that each of the cyber assets are needed for, and the mission-critical operations that each of those tasks supports. A visual scene like this, with three layers consisting of one grid and two rings of spheres can be used to represent many different types of data. For example, the middle ring, which is used in Figure 2 to represent mission-critical tasks, could be used in another scene to represent security events that have occurred on the cyber assets below it. In that case, the spheres in the middle ring could be color-coded to represent the severity of the security events, and the spheres in the top ring could be color-coded red, yellow or green to represent status of the mission-critical operation as a result of the security event.

Figure 3 depicts layered grids. The bottom grid is comprised of only those assets that are needed for mission planning tasks. The SecureScope visualization system, on which these scenes were built, allows the user to constrain the type of data retrieved from the database. In this example the user asked to only see IA data about assets needed for mission planning. However, the

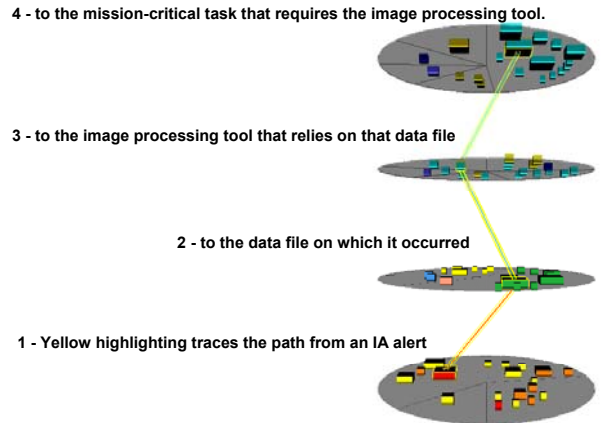
analyst could have constrained the query in many other ways, such as by time, location, phase of the mission, or protocol. The top grid is comprised of IA alerts that have been reported on the cyber assets needed for mission planning. The alerts are colored to reflect their severity. The user could add more layers, if desired, up to a total of four.

In the future we are planning to add a map layer so that the assets or the alerts could be associated with geographic points. The next developmental step will be to add a network topology layer so that the assets or alerts could be related to where they are in the network topology.

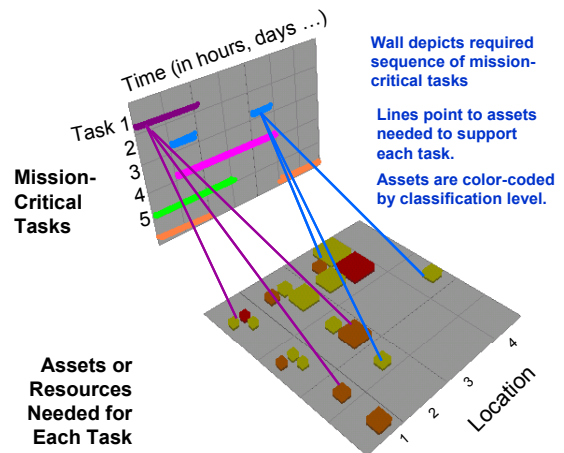


**Figure 3 – A layered grid scene showing the assets needed for mission planning that have some type of IA event (e.g. probe, suspicious activity) on them.**

Figure 4 shows a layered wheel scene. Each wheel is subdivided into sections, like a pie chart. The lowest layer is comprised of IA alerts. They are sectioned into groups of alerts and positioned on the wheel based on these groupings. The layer above the alerts is used to depict cyber assets that are singular in nature, such as a single data file, workstation, or software application. The next layer up includes complex cyber assets or capabilities that are comprised of singular assets or capabilities. For example, PowerPoint could be considered a singular asset and placed on the layer that is second from the bottom, while Microsoft Office would be placed on the next layer up with the more complex assets or capabilities. The top layer is used to depict mission-critical tasks that rely on the cyber assets below it. From this type of scene a user can trace the impact of a security event from the bottom layer up to the top, to find out which tasks will be affected by the alert. The example in Figure 4 can be used for the classic “bottom-up” mission impact analysis.



**Figure 4 – A layered wheel scene that traces mission impact from the security alert, to the assets affected by the alert, to the mission-critical task affected by the breached assets.**



**Figure 5 – Mission resource wall shows the cyber assets that are needed for each mission-critical task, the estimated time they are needed, and the order in which they are needed.**

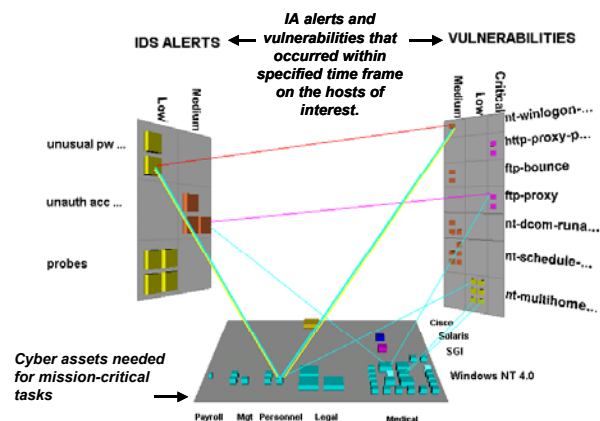
In Figure 5 we capture time and sequence elements of mission-critical tasks and relate those to the mission impact of security events. The vertical wall displays information about: 1) time, for which the user specifies the range (e.g. January 1-10, 2000) and granularity (e.g. days, hours, seconds); and 2) mission-critical tasks. The duration bars represent the estimated time the task will be performed. The horizontal “floor” in the foreground is a grid populated with boxes that represent the assets or capabilities required of the tasks depicted on the wall.

The boxes are positioned on the grid based on two x/y properties, such as the location of the assets and the organization that owns those assets. The boxes may be colored, sized, set to blink, or put into motion to communicate additional information about the assets. For example, the boxes may be color-coded to indicate their security classification level, and size-coded to indicate their general level of criticality to the organization that owns them. The association lines show the specific time interval that each asset is needed for specific mission-critical tasks.

Other versions of this mission resource wall contain a rear grid, in addition to the front grid. The rear grid can be used to display information about alerts that have occurred on the critical assets. Using such a scene the analyst can see the assets needed for each mission-critical task, the security alerts that are occurring on those assets, and which tasks may be affected by those alerts.

Figure 6 is a conceptual design that has not been fully developed and tested at the time of this writing. It consists of several surfaces that can contain information about missions, assets, alerts, time, etc. In the example in Figure 6, there are several mission-critical network-based assets depicted as boxes on the grid. They are positioned on the grid based on the organization that owns them and the operating system they are running. The vulnerabilities that were detected on those assets during the last regular scan are shown on the right wall. The IDS alerts that have been reported on those assets are on the left wall. Lines emanating from the grid to the walls associated specific network assets with the vulnerabilities and alerts detected on them. The lines running between the vulnerabilities wall and the alerts wall represent a linking between the known vulnerabilities and apparent exploits of those specific vulnerabilities.

This room-like view of the data relevant to mission impact can be modified to include additional surfaces. Up to four walls and a “ceiling” grid can be added to the scene. One of the walls could be programmed to depict time, in a manner similar to that shown in the previous Figure 5. Another wall or grid could be used to represent a geographic map that can show where alerts on the wall have occurred.



**Figure 6 - A room-like view of mission-relevant security data is captured in this scene, with which the analyst can relate known vulnerabilities on critical assets to security alerts that have already occurred on those assets.**

### 3. Technology Underlying the Mission Impact Visualizations

#### 3.1 Data Populations

The 3D visualizations produced by Secure Decisions’ system are populated from data contained in a relational database. Multiple dimensions of a data population, or table, can be viewed simultaneously, as well as relationships with other data populations. Thus, the user can use the 3D visualization to see relationships within a single data population or between multiple data populations.


Each table in a database schema represents a single data population; fields within that table represent various dimensions of that data population. For example, a table of security events represents the data population called Events. Fields within the Events table categorize the events as vulnerabilities, suspicious activities, or attacks. Other fields within the Events table identify the severity of its items as high, medium or low. A 3-D visualization of the data population called Events could visually correlate the severity of events with their category names. A different table, named Host table, could contain information about a population of workstations on the network. The Hosts data population could be further described by fields such as: Host name, Host operating system, or Host location. A 3-D visualization could be used to show how the two data populations

called Events and Hosts relate to each other, and how the various dimensions of each population (e.g. severity of Event, Host OS) interrelate.

Secure Decisions' visualization system does not analyze the data itself or suggest interesting ways to view the data. Rather, the visualization system provides model-driven analysis that allows for a non-SQL savvy end-user to configure and populate 3D, intuitive data models

### 3.3 Architecture

The visualization system conforms to the client/server model, as depicted in Figure 7 below. The use of a red font in the figure distinguishes those portions that belong to the visualization system from those that an end-user would supply. The end-user configures the visualizations via the Console (the client), which sends the high-level data request via Java's Remote Method Invocation to the Data Service (the server), which then interprets and translates it into one or more SQL statements, which are submitted to the database and executed via Java JDBC. The data is then returned to the Data Service, translated, and returned to the Console. The previously selected scene type is then populated with data. As data is updated or inserted into the database schema, the visualization does not dynamically update the scene. Instead, the scene is updated only after the point at which the user makes an explicit request to refresh the scene. At that time, the data, obeying any restrictive criteria accompanying the request, is fetched from the database and returned to the Console.



**Figure 7 – SecureScope architecture underlying mission impact visualizations**

The Console was built using a combination of technologies, including C++, Java, VRML and COM. The Data Service is 100% Java. The visualization can connect to most relational database schemas running in Oracle, Microsoft Access, Microsoft SQL Server and

MySQL relational database management products. Barring support for JDBC and basic SQL capabilities, nothing about the design of the visualization system would prevent it from working with other database systems, however, further development work would be necessary, as there are typically differences in SQL syntax from one database product to another.

### 3.4 Database Schema

The mission impact visualizations were designed to interface with most standard mission impact schema. However, in the absence of an existing schema, we can work with the analyst to develop an appropriate schema.

## 4. References

- [1] D'Amico, Anita and Larkin, M., "Methods of Visualizing Temporal Patterns in and Mission Impact of Computer Security Breaches", *DISCEX II*, Anaheim, CA, June 2001
- [2] Tufte, Edward R., *Envisioning Information*, Graphics Press, Cheshire Connecticut, 1990.