# Models of Information Security Trend Analysis

Tim Shimeall, Ph.D., Phil Williams, Ph.D.
CERT® Analysis Center, Software Engineering Institute
Carnegie Mellon University, Pittsburgh, PA

**ABSTRACT**

This paper discusses a framework for conduct of information security trend analyses. While several organizations are performing such analyses, there is wide disparity between the level of the analyses, the applicability of results, and the assumptions involved in properly interpreting the results. The framework offers a common ground in which these issues may be resolved. An example analysis process is presented in the paper. The paper includes a discussion of cautionary factors in the application of this framework.
Keywords: Information Security, Trend Analysis, Network Security, Law Enforcement, Electronic Crime

## 1. Introduction

An information security incident is an event that negatively impacts the security of one or more sites. Network security events are growing, both in frequency and in impact, but quantitative data regarding security incidents is both difficult to obtain and hard to place into a meaningful framework. While several organizations are performing such analyses, there is wide disparity between the level of the analyses, the applicability of results, and the assumptions involved in properly interpreting the results. This paper reports on models for trend analysis, developed at the CERT Analysis Center to provide such a meaningful framework.

Merriam Webster's defines "trend" as "(1) to extend in a general direction: follow a general course or veer in a new direction and (2) to show a tendency for example, to incline or trend upwards or to become deflected or shift". Following this, trend analysis can be understood as a search for patterns over time in order to identify the ways in which they change and develop, veer in new directions or shift. Not surprisingly, trend analysis has become ubiquitous. A simple Internet search for trend analysis, for example, generates over 1.4 million hits. This reflects the fact that trend analysis is used in an enormous variety of fields from the analysis of global politics and economic and social issues, to legal assessments, engineering, and demographics. Trend analysis is used in efforts to understand and predict financial markets. It offers a dynamic assessment of the past and present as well as a means of extrapolating into the future. Moreover, trend analysis can take different forms. It can be both quantitative and qualitative in nature, providing a basis for precise assessments of change as well as for broad assessments of the direction and impact of change. Where trend analysis is done well, it can assist in identifying many things including the direction and scope of the changes taking place, whether they are positive or negative in their impact and whether they are qualitative or quantitative in nature. In some cases this can provide a basis for anticipating future behavior and – in cases where this behavior is damaging – taking precautionary measures to defend against it. As the Institute for Alternative Futures noted: "the first step in thinking about the future involves exploring trends that are already underway. A trend is a pattern of change over time in things of importance to the observer…. for any area it is important to identify the key trends, or patterns of change, shaping it…. The earlier a trend is detected, the greater is the flexibility of organizations responding to it. Trends often become issues or even crises before measures are taken. At the crisis stage the decision costs for organizations are usually lowered but the range of options is narrowed."[10] The implication of this is that understanding trends is an important tool in the early detection of problems and challenges. Indeed, in the area of information security, enhanced understanding of trends, patterns, and anomalies could contribute significantly to indicators and warning processes that are a key component of efforts to anticipate, thwart, or mitigate intrusions. It is possible, for example, to extrapolate trends so that defenders have at least some expectation about broad developments that might occur. While this is not foolproof by any means, it can provide some basis for anticipation and lessen surprises.

® CERT and CERT Coordination Center are registered with the US Trademark and Patent Office.

## 2.Types of Trends

Before considering the types of trends we want to identify it is important to acknowledge the relationship between trend analysis and pattern analysis. Indeed in analyzing trends we are, in effect, analyzing changing patterns of activity over time – and when we suggest that there is a trend in activity we are merely saying that the dominant pattern of activity is changing over time. There are various forms this change can take. It can be gradual or abrupt, it can include regularities as well as anomalies or deviations from established patterns, it can have continuities as well as gradual or abrupt departures.

### *2.1. Internal and external patterns*

Within a trend analysis, therefore, it is important to look for patterns in three different meanings of the term: (1) recurring themes and motifs that are common across a whole range of cases at different times and places; (2) a particular design or a particular distribution of events that is either repeated in regular ways or represents an anomaly or deviation; (3) a dominant form of activity that shifts over time and is replaced by another dominant form. It is also useful to distinguish between internal patterns that are, in effect, intrinsic to the activity, and external patterns that increasingly characterize a large portion of the activity. Internal patterns may provide insight into the modus operendi (MO) of an intruder, providing law enforcement with a basis for identification or characterization. External patterns may facilitate strategies in facing larger issues in cyber crime. In addition, it is useful to see to what extent the following kinds of trends are evident in the dataset.

### *2.2. Temporal trends*

The key question here is "what does the analysis reveal about the temporal nature of intrusions? In answering this, the analysis seeks to identify the particular distribution over time of the events in the data-set. It looks for regularities within the distribution patterns. Are there, for example, cyclical patterns within the data? Are there seasonal patterns? Do we see recurring patterns of peaks and valleys or of more gradual increases and decreases in the frequency of the phenomenon being examined? What kinds of shifts do we see in levels of intensity and frequency? Are they related to particular triggers? In this connection, there are three dimensions that can be considered: 1) the extent to which different instances of the behavior (or cases) are clustered together in a time period, 2)the intensity of the activity (i.e., how many instances are there of this behavior in the period identified), and 3) the duration of the activity. This is important because the cessation of an activity can itself be meaningful. It might signify, for example, that a particular objective has been achieved or that those engaged in the activity have moved on to other things. Intruders frequently are passionately impressed with their own anonymity, and may cease activity if they perceive a chance of being discovered. Defenders and law enforcement agents may seek to produce conditions that promote such perceptions, as a preventative strategy.

### *2.3. Spatial trends*

The idea of spatial trends in cyber-attacks superficially appears rather strange. Yet the targets of computer intrusions exist in physical space as do the intruders and their tools. In thinking about spatial trends, the relationship between physical space and territory and cyber-space may be a critical determinant of the categorizations that are typically made in discussing intrusions. The notion of cyber-warfare, for example, presumes that one nation-state is attacking computer and information systems that operate in another nation-state. In thinking about spatial trends in cyber-space it is also possible to look at sectors. Are particular sectors (e.g. schools and universities) a major source of intruder behavior? Similarly, are major economic sectors such as banks and other financial institutions a major target? These issues may facilitate prioritizing law enforcement efforts.

### *2.4. Associational trends*

Another critical question about computer intrusions is the extent to which particular incidents are linked together – commonalities, for example, in the targets of intruders, in the methods that are used, or in the timing of attacks might suggest that there are growing linkages amongst members of the intruder community that include not only a willingness to share tools but also to coordinate their actions. John Arquilla and David Ronfeldt [1] have examined the notion of swarming as an attack strategy that can be

used in both physical space and cyber-space. A close examination of associational trends might provide evidence that this kind of action is taking place – or not. The critical point here is that it is important to consider the possibility that incidents might be linked in certain ways. These linkages may also provide insight into whether the intrusion is an end in itself, or part of a larger electronic crime.

### *2.5. Compound trends*

It is also possible to look for compound or complex patterns within the trend data. These are the kinds of patterns that combine several different dimensions such as time and space. In effect, what we are doing here is identifying two separate axes (temporal and spatial) and considering what new patterns might be detected as a result of considering them together rather than as separate and independent from one another.

In other words, trend analysis of computer intrusion data can help to illuminate aspects of the intruder community and the diffusion and use of intruder tools, commonalities, associations, or links among incidents that had hitherto been regarded as separate and distinct from each other; the importance of possible triggers ranging from reports of new vulnerabilities to world events such as increased tensions in the Middle East or between China and the United States. In effect, trend analysis can help provide provisional answers to the who what where, when and how (and to a lesser extent the why) questions relating to intruder behavior. As any investigator knows, such questions are critical to building an effective prosecution. Before looking more closely at the various cyber trends, however, it is important to acknowledge the limits and problems of trend analysis.

## 3. The Limits of Trend Analysis

Trend analysis of computer intrusions is essentially about making inferences from data that is a partial, incomplete and possibly unrepresentative sample of the phenomenon under investigation. In effect, it offers a series of narrow snapshots of a much broader phenomenon. Moreover, the reports on which the analysis is based usually provide only the immediate data on the intrusion. There is often no follow-up or, if there is, it is not reported. Consequently, the data suffers not only because of possible sampling biases (especially since all reports are voluntary) but also because it represents a snapshot in what is actually a moving and dynamic picture. In addition, the requirements of confidentiality make it difficult to undertake particular kinds of analysis such as victim profiling. It is possible, for example, that intruders are systematically targeting a set of financial institutions. Even if all of them reported this to CERT/CC, it would not necessarily be reflected in the data as the identity of victims of intrusions is confidential. The other dimension of what might be termed the input or evidence problem is that computer intrusions represent a form of activity in which every effort is made to maintain anonymity and obfuscate any trail back to the perpetrators.

A second set of problems concerns the difficulty of trend analysis as such. One of the difficulties, for example, concerns the baseline. Without a baseline, it is easy to draw the wrong conclusions. Yet determining an appropriate baseline in an environment that is highly dynamic is very problematic. The number of incidents, for example, has gone up very dramatically, but so has the number of computers in operation and therefore the number of available targets, the number of possible attack vectors, and the number of computers from which attacks can be launched. In fact, the Internet is growing faster than the rate of intrusions on the Internet. There is a danger with focusing on one set of trends – and drawing what are in effect internal conclusions about them when there are other broader or external trends that might be having a major impact.

A third set of problems center on interpretation. It is all too easy, for example, to slip from correlation to causation. A spike in the number of intrusions that accompany an international crisis could be a direct result of that crisis. Yet it is also possible that what are seen as directly connected actions are not in fact linked. A DDoS attack that coincided with an international crisis, for example, might well have been in preparation and underway for some time – and have absolutely nothing to do with the crisis. In other words it is important to beware of what can be termed the coincidence problem.

The implication of all this is two-fold. First it is necessary to be extremely cautious in the kinds of conclusions that are drawn from the trend analysis that is being undertaken here. In effect, simply being

sensitive to these problems makes it less likely that inappropriate inferences will be drawn or that there will be excessive confidence in the internal findings. Second, it is useful to consider strategies for dealing with trend analysis when it goes wrong. In some cases, for example, it is necessary to go from trend analysis focusing strictly on the dataset to a broader set of data that permits some kind of fusion analysis in which the tentative conclusions that have been drawn can be considered against a broader knowledge base.

To acknowledge the limitations of trend analysis of computer intrusions, however, does not diminish the very real benefits that can be obtained from the systematic and thorough analysis of the data-set. As the next section suggests, trend analysis of the CERT Coordination Center (CERT/CC) data set can yield some very illuminating insights.

## 4.Cyber Trends

Trends of Internet security incidents exemplify each of the types described in the previous section. The data in this section is taken from reports to the CERT/CC. The CERT/CC is a center of Internet security expertise at the Software Engineering Institute, a federally funded research and development center operated by Carnegie Mellon University. The CERT/CC receives reports of security incidents from voluntary reporters worldwide, and provides assistance in dealing with vulnerabilities, handling incidents, and other security activities. The reports have more than doubled each year since 1998, growing from 3,734 incidents in 1998 to 52,658 in 2001.[2] From this large pool of data, the CERT Analysis Center has been conducting statistical and in-depth studies of incidents with substantive impact on the victim organizations. These studies have provide examples of incident trend analysis.

An internal pattern in information security is a series of events that flow together to constitute a security incident. For example, there have been several incidents reported to CERT/CC where the intruder has worked in distinct stages. See Figure 1 for a diagram describing one example. In this incident, the intruder initially conducted a series of probes to locate easily-compromised computers. He or she then selected three computers, illicitly gained administrative access on them, and installed attack software on each of these computers (labeled 1, 2 and 3 in Figure 1). When later triggered, the attack software caused each of these computers to contact the target site (shown on the upper right in Figure 1) in a carefully-controlled manner, which cumulatively caused a compromise at the target. The internal pattern of distribution of attack is shown in this example, but is also present in distributed denial of service attacks (where a group of compromised computers act to consume available network resources at a set of targets) and distributed sensing attacks (where a group of compromised computers act to intercept traffic intended for a set of targets). In each of these cases, the flow of activities from probe through compromise is exemplified. From a law enforcement perspective, awareness of internal patterns aid in understanding how closely an individual may be tied to a style of intrusion. This may be essential in identifying a perpetrator.

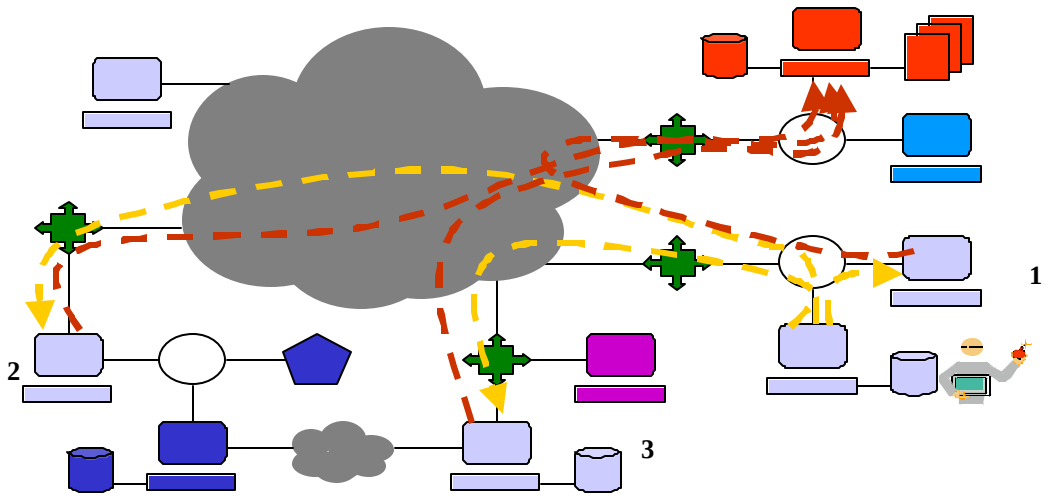


Figure 1: Staged Incident

An external pattern in information security is a flow among a series of incidents that may occur at a common site or set of sites. The flow here is primarily a flow of knowledge or information that develops intrusion behavior. An example of this is a preliminary scan for network vulnerabilities by a set of intruders, followed later by a compromise by different sets of intruders. The implication is that intruders share information on vulnerable sites, and use information gathered by others. There are statistical probabilities that document such trends [9]. This may provide investigators the opportunity to identify perpetrators through subsequent intrusions. Other examples are a series of contacts that develop a set of targets for later intrusion, or characteristics for later tool development. As intruder tools are detected and analyzed, there have been documented cases of developers using the analysis as a basis for more sophisticated intruder tools (as is diagrammed in Figure 2). There have been documented cases where intruder groups compete with each other to produce more and more harmful tools. In Figure 2, the dark boxes represent attack tools successively developed by one of two groups of intruders. The light boxes represent analyses of such tools by security experts, often documenting the detectable characteristics and preventative measures. As shown in the figure, the competing intruders use both previously developed code (both from their tools and their competitor's) and analyses of this code to produce later tools.

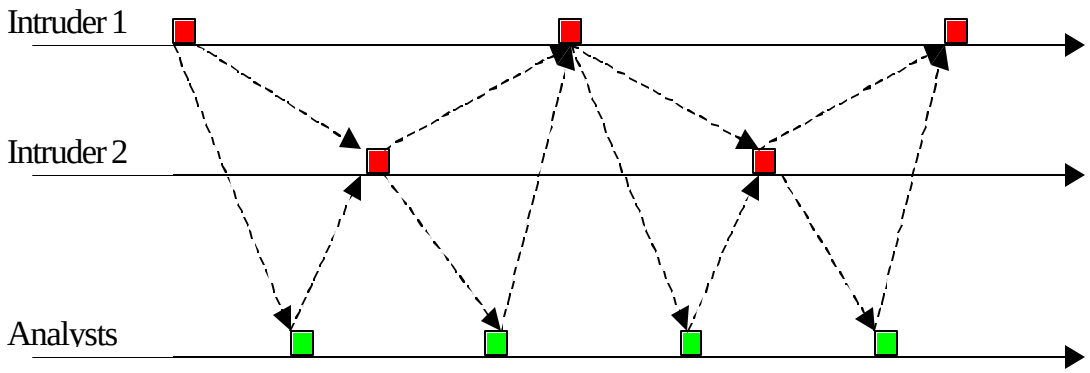


Figure 2: Flow of Information in Competing Tool Development and Analysis

A temporal trend is one in which the relationship between incidents and real time is significant. One type of such trend relates to the timing between an event that may trigger an incident (say, a new product announcement, a provoking statement, or the release of a new intrusion tool) and the corresponding incidents. This trend provides insight on whether the interval between stimuli and response is increasing, decreasing, becoming more consistent, or becoming more varied. Another type of temporal trend examines the times at which incidents occur. This trend provides insight on whether events are periodic or aperiodic (particularly untriggered aperiodic). One example of stimulus-response trending is shown in Figure 3.

The black line in Figure 3 shows the number of incidents with serious effects on a reporting site as reported to the CERT/CC each week between June 24, 2000 and Feb 17, 2001. In this graph and ones that follow, the incident counts shown exclude simple port-scanning and other incidents with little or no lasting impact on the targeted sites (e.g., failed fraud, unexploited vulnerabilities, false alarms, hoaxes). The gray arrows are major announced hacking conventions (e.g., Defcon, Black Hat Briefings, etc.) Those conventions that were based in the United States appeared at or close to local peaks in the incident reporting. The gray vertical lines show the number of new or revised exploits published on a full-disclosure website(www.packetstormsecurity.org). Peaks in the exploit publication rate were weakly correlated (at around a 80 percent confidence level) with peaks in the incident reporting rate one to three weeks later. In both cases, activity that equipped the intruder community could be associated with increases in incident reporting. While the incident reporting curve is not completely explained by these stimuli, a partial trend may be discerned.

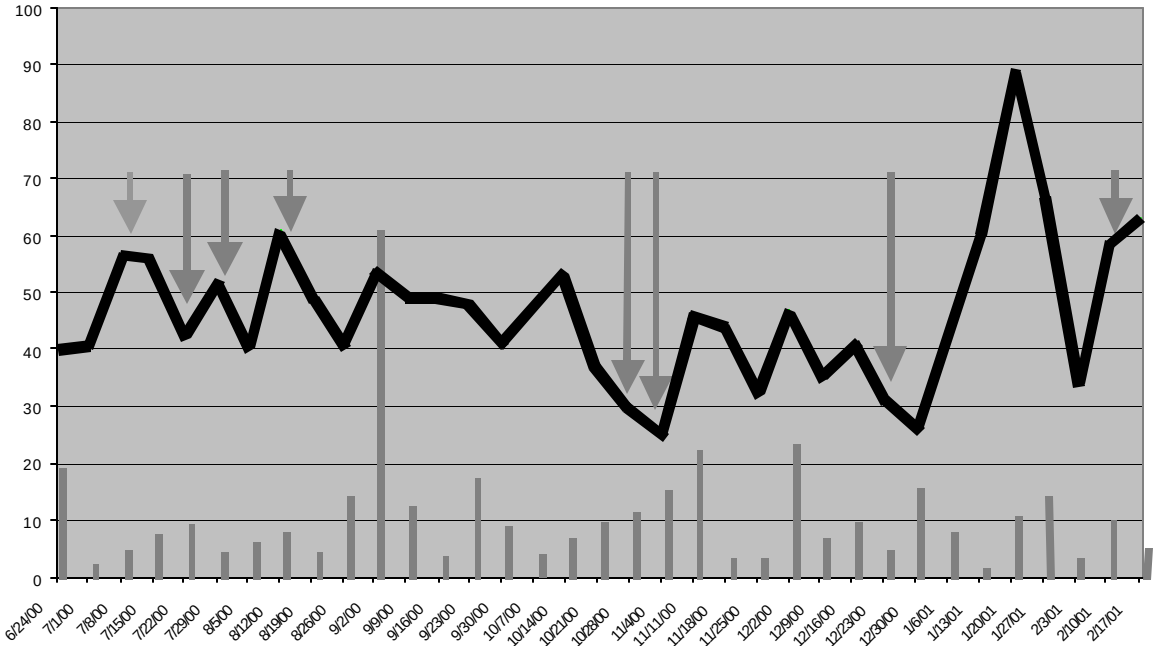


Figure 3: Example of Stimulus-Response Trending

Figure 4 diagrams a series of trends related to the exploitation of vulnerabilities in reported security incidents. The vertical axis in Figure 4 is expanded in comparison to Figure 3, reflecting that many reported security incidents lack sufficient detail to allow identification of the vulnerability being exploited. In Figure 4, the most commonly-exploited vulnerability in incidents were the WU-FTP (format string)[3] and RPC (statd overflow) [4] vulnerabilities. These reflect long-term exploitation of vulnerabilities in widely-used network services. There are also shorter-term exploitation patterns, such as for a vulnerability in telnet [6] where the activity started in mid-August, and dropped off after early December, and a vulnerability in LP [5] which had some initial probing in the August-September timeframe prior to larger scale exploitation in January-February. Overlapping cycles of vulnerability exploitation are common. Network administrators cannot depend on a "top ten" list – the intruders tend to modify their tactics too quickly. It is also true that while substantive periods of time may elapse between the discovery of a vulnerability and its widespread exploitation, there is a trend toward more rapid exploitation.

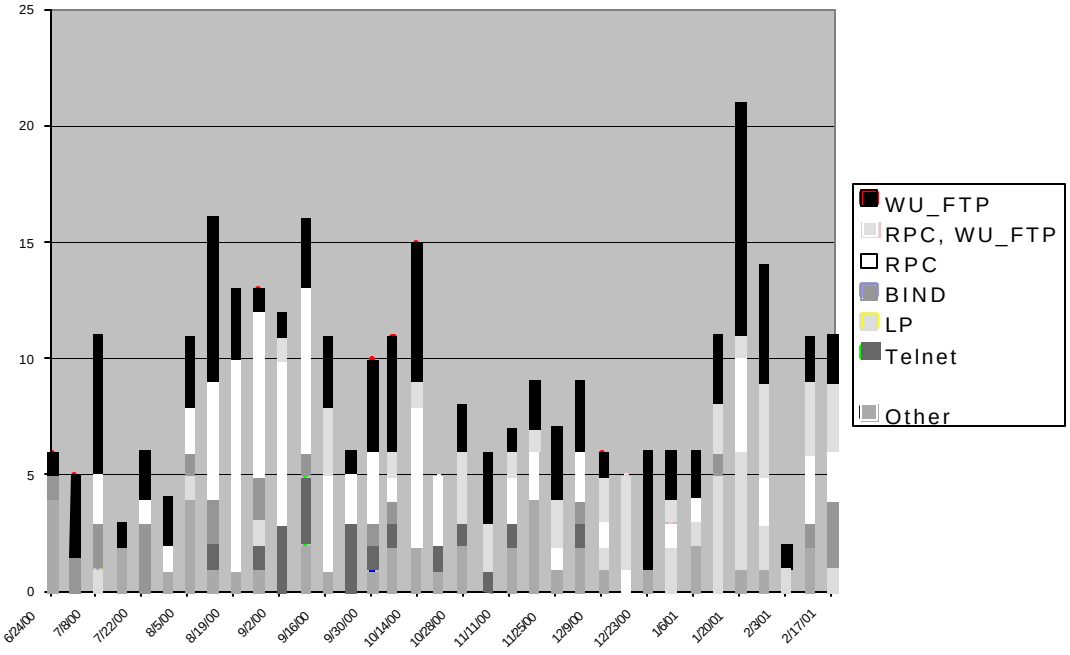


Figure 4: Vulnerabilities associated with Reported Incidents

A spatial trend is one in which the physical or logical location of incidents is considered. An intruder may choose to concentrate on one particular local area or a restricted range of IP addresses. This will likely indicate that the intruder has some rationale behind this localization – and this may help an administrator determine how to defend against such an intruder. Attack tools have also been identified that focus on specific regions. CodeRedII [7] acted differently on Asian networks than on occidental ones. Examining incidents through spatial trending may provide insight into how intruders target their attacks.

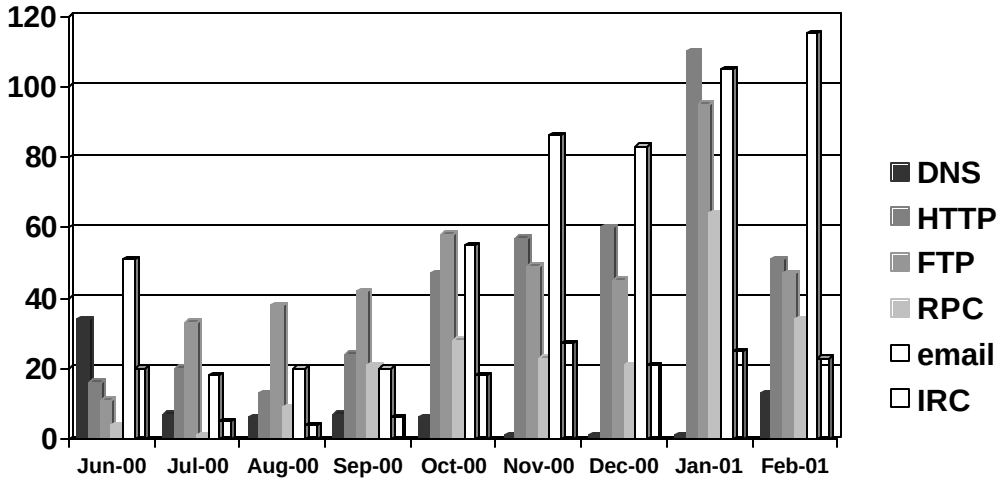


Figure 5: Services in Incidents

An associational trend is one that seeks to discover the commonalities between incidents. On the surface, the incidents may appear to differ greatly, but through analyzing for common factors, a trend may emerge. One example is shown in Figure 5, which shows use of services in incidents (which includes use of the service as an attack vector and threat sources as well as vulnerability to exploit). By trending the service use in incidents, the analyst may identify which are the most common attack vectors and exploited services, not just as a static picture but as a dynamic environment of intrusion. In Figure 5, the period of June-July 2000 saw a rapid tradeoff of intrusions involving DNS and FTP. IRC and email show also significant variations in numbers of intrusions. In contrast, the numbers of incidents involving HTTP remain relatively constant over the period. Another example of an associational trend would be a series of web defacements with common elements in the messages.

A compound trend combines several of the other trends to provide a more complete picture. An example compound trend might be incidents that form both temporal and spatial trends. Incidents strongly linked to a similar time and location may indicate a focused series of intrusions, and provide a useful basis for merging non-intrusion information for analysis of causes. Compound trends seek to find increased significance in the combination of other trends. Rather than looking at one dimension of variation, compound trends seek to explore several dimensions at once. Another example might be examining the exploitation of vulnerabilities in combination with the number of hosts involved in incidents. This would lead to some understanding of which vulnerability exploits tend to produces rapidly-spreading incidents, as opposed to focused and relatively local incidents. Over time, the study of compound trends may give insight into productive strategies in network intrusion. Currently, the focus of effort is very tactical in nature: how to exploit specific vulnerabilities or defend against such exploits. For more proactive defense, understanding of network defensive and offensive strategies must be developed. Such development requires a basis of data, and compound trending offers the opportunity to acquire the required data in a useful form.

In the process of trend analysis, it is often illustrative to discover not just the trends, but also deviations from trends. A clear and recurring pattern that appears to be well-established provides a baseline from

which it is possible to detect and assess changes, departures, or deviations. By examining deviations from an observed baseline, the analyst may discover false or coincidental trending and produce a deeper and more accurate analysis. Probably the most famous example of this is Kepler's attention to the deviation of the orbit of Mars from his spherical model of the orbit of planets, which lead to his discovery of the laws of planetary motion. In incident activity, deviations from trends should be treated with interest. Such changes either provide indicators of an anomaly that needs to be further examined or of something that is recognizable and understood as a distinct trend – a shift from one pattern of behavior to another. As an example, consider Figure 5, as discussed above. If the expected trend is for the number of incidents involving a particular service to grow, plateau and decay over time (as is shown with DNS, IRC and FTP), why do email and HTTP display relatively constant numbers? It may be that the period of observation is too short to show the growth and decay in incident numbers for these services. It may also be that since HTTP and email are widely accepted on the Internet, they form a continual vehicle for incidents (although such an analysis would need to recognize that FTP and DNS are also almost-universally supported services). A third explanation could be that how these two services are used in incidents varies more than the other services. Email and HTTP may serve as attack vectors, as means of responding to attack, or as notification of attack (although the same might be said of IRC). At this time, the best that may be said is that the deviation of these two services from the more general trend is unexplained – which is a status that offers the opportunity for interesting future insights.

## 5. Process of trending

A modern computer network is an exceedingly dynamic environment. New hosts, services, and users are constantly being added. New attacks and failures are constantly arising. Developing an analysis process that will deal with this environment and produce meaningful trends for security incidents is a challenging task. This paper describes one such analytical model, developed at the CERT Analysis Center, building on processes described previously [8].

The first aspect of incident trend analysis is the establishment of an appropriate flow of incident information. This is not a trivial task, as reporting organizations may be slow to establish trust relationships with the analytical organization, and until such trust relations are established, incident reports are likely to be cursory and incomplete, at best. The more diverse the reporting base contributing to the flow, the more robust the trend analysis results are likely to be.

The next aspect is to identify appropriate profiles for the reported incidents. Incident reports may be expected to vary greatly in terms of detail. Partly, the variation is because security incidents, by the action of the intruder, may leave incomplete records of the malicious actions taken. Partly, the variation is because system administrators often have very little available time for reporting of security incidents. By identifying an appropriate profile, the analyst may be able to normalize this variation in data. In addition, the profile will help to isolate variables of interest for trend generation and analysis. Finally, the profile forms a basis for consistent comparisons and statistical distributions. Once the profile is identified, each reported incident should be profiled for analysis.

Once the incidents are profiled, the next task is to locate and correlate meaningful outside data sources for analysis. The analyst should strive for a diverse set of data sources, to provide as much insight as possible for the readers. The effort should be to focus on relevancy of data to profiled incidents, within an overarching indications and warnings process. Relevancy should be identified in terms of timing, causation and implication arguments.

Finally, as trends are developed, gaps in the available insights will become evident. The trend generation process will then iterate, by collecting and correlating further information and by improvements in the incident profiles. The process can be understood in the following steps: (1) The establishment of baseline patterns as a trend that is understood and regularly monitored. (2) The detection of a change in the pattern – an anomaly or deviation from the norm. (3) Determination that the anomaly is : either indicative of change to a new trend that is familiar and understood as such (pattern recognition); or a blip or aberration that is not indicative of a new trend (true anomaly); or a development, the significance of which is uncertain, but

that requires further scrutiny because of the possibility that it represents a new trend (uncertain anomaly); or indicative of a change to a new trend the purpose, meaning and significance of which is not yet understood (pattern discovery).

## 6. Conclusion

This paper has described models of information security trend analysis. By typifying trends and associated patterns, these become more accessible tools. The types provide a basis for clarifying the goals of the analysis. Applied to information security, the relative importance of distinct aspects of the subject matter provides a basis for selection of trends to analyze. For example, a project manager concerned about the schedule of access to a new network service may consider temporal trends to be most important. Law enforcement agents studying patterns of computer crime may use associational trending.

As on-line crime increases and societal perception of the importance of the Internet grows, understanding trends in security incidents will become more important. The CERT Analysis Center is continuing to work in this area, exploring methods for analysis and application of the insights to infrastructure protection and law enforcement issues.

## 7. References


1. John Arquilla and David Ronfeldt, Swarming and the Future of Conflict, RAND Report DB-311-OSD, 2000.

2. CERT Coordination Center, “CERT/CC Statistics 1988-2001”, Software Engineering Institute, Carnegie Mellon University, available at http://www.cert.org/stats/cert_stats.html.

3. CERT Coordination Center, “CERT® Advisory CA-2000-13 Two Input Validation Problems In FTPD”, Software Engineering Institute, Carnegie Mellon University, available at http://www.cert.org/advisories/CA-2000-13.html.

4. CERT Coordination Center, “CERT® Advisory CA-2000-17 Input Validation Problem in rpc.statd”, Software Engineering Institute, Carnegie Mellon University, available at http://www.cert.org/advisories/CA-2000-17.html.

5. CERT Coordination Center, “CERT® Advisory CA-2000-22 Input Validation Problems in LPRng”, Software Engineering Institute, Carnegie Mellon University, available at http://www.cert.org/advisories/CA-2000-22.html .

6. CERT Coordination Center, “CERT® Advisory CA-2001-21 Buffer Overflow in telnetd”, Software Engineering Institute, Carnegie Mellon University, available at http://www.cert.org/advisories/CA-2001-21.html.

7. CERT Coordination Center, “CERT® Advisory CA-2001-23 Continued Threat of the ‘Code Red’ Worm”, Software Engineering Institute, Carnegie Mellon University, available at http://www.cert.org/advisories/CA-2001-23.html.

8. Casey Dunlevy, Timothy Shimeall, and Phil Williams, "Intelligence Analysis for Internet Security: Ideas, Barriers and Possibilities", SPIE Law Enforcement Conference, Boston MA, October 2000.

9. Suomo Moitra, and Suresh Konda, “A Simulation Model for Managing Survivability of Networked Information Systems”, Submitted for Publication, 1999.

10. Institute for Alternative Futures, “Environmental Scans and Trend Analysis”, available at http://www.altfutures.com/svcs/enviroscans.htm.