Internet Society's *Networks and Distributed Systems Security Symposium*, March 1998.

Live Traffic Analysis of TCP/IP Gateways

Phillip A. Porras
porras@csl.sri.com
Computer Science Laboratory

SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025

Alfonso Valdes
avaldes@csl.sri.com
Electromagnetic and Remote
Sensing Laboratory
SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025

The work presented in this paper is currently funded by
DARPA/ITO under contract number F30602-96-C-0294.

Point of Contact: Phillip A. Porras
Phone: (415) 859-3232
Fax: (415) 859-2844

November 10 1997

ABSTRACT

We enumerate a variety of ways to extend both statistical and signature-based intrusion-detection analysis techniques to monitor network traffic. Specifically, we present techniques to analyze TCP/IP packet streams that flow through network gateways for signs of malicious activity, nonmalicious failures, and other exceptional events. The intent is to demonstrate, by example, the utility of introducing gateway surveillance mechanisms to monitor network traffic. We present this discussion of gateway surveillance mechanisms as complementary to the filtering mechanisms of a large enterprise network, and illustrate the usefulness of surveillance in directly enhancing the security and stability of network operations.

1. Introduction

Mechanisms for parsing and filtering hostile external network traffic [2],[4] that could reach internal network services have become widely accepted as prerequisites for limiting the exposure of internal network assets while maintaining interconnectivity with external networks. The encoding of filtering rules for packet- or transport-layer communication should be enforced at entry points between internal networks and external traffic. Developing filtering rules that strike an optimal balance between the restrictiveness necessary to suppress the entry of unwanted traffic, while allowing the necessary flows demanded for user functionality, can be a nontrivial exercise [3].

In addition to intelligent filtering, there have been various developments in recent years in passive surveillance mechanisms to monitor network traffic for signs of malicious or anomalous (e.g., potentially erroneous) activity. Such tools attempt to provide network administrators timely insight into noteworthy exceptional activity. Real-time monitoring promises an added dimension of control and insight into the flow of traffic between the internal network and its external environment. The insight gained through fielded network traffic monitors could also aid sites in enhancing the effectiveness of their firewall filtering rules.

However, traffic monitoring is not a free activity--especially live traffic monitoring. In presenting our discussion of network analysis techniques, we fully realize the costs they imply with respect to computational resources and human oversight. For example, obtaining the necessary input for surveillance involves the deployment of instrumentation to parse, filter, and format event streams derived from potentially high-volume packet transmissions. Complex event analysis, response logic, and human management of the analysis units also introduce costs. Clearly, the introduction of network surveillance mechanisms on top of already-deployed protective traffic filters is an expense that requires justification. In this paper, we outline the benefits of our techniques and seek to persuade the reader that the costs can be worthwhile.

2. Toward Generalized Network Surveillance

The techniques presented in this paper are extensions of earlier work by SRI in developing analytical methods for detecting anomalous or known intrusive activity [1], [5], [12], [13]. Our earlier intrusion-detection efforts in developing IDES (Intrusion Detection Expert System) and later NIDES (Next-Generation Intrusion Detection Expert System) were oriented toward the surveillance of; user-session and host-layer activity. This previous focus on session activity within host boundaries is understandable given that the primary input to intrusion-detection tools, audit data, is produced by mechanisms that tend to be locally administered within a single host or domain. However, as the importance of network security has grown, so too has the need to expand intrusion-detection technology to address network infrastructure and services. In our current research effort, EMERALD (Event Monitoring Enabling Responses to Anomalous Live Disturbances), we explore the extension of our intrusion-detection methods to the analysis of network activity.

Network monitoring, in the context of fault detection and diagnosis for computer network and telecommunication environments, has been studied extensively by the network management and alarm correlation community [8], [11], [15], [16]. The high-volume distributed event correlation technology promoted in some projects provides an excellent foundation for building truly scalable network-aware surveillance technology for misuse. However, these efforts focus primarily on the health and status (fault detection and/or diagnosis) or performance of the target network, and do not cover the detection of intentionally abusive traffic. Indeed, some simplifications in the fault analysis and diagnosis community (e.g., assumptions of stateless correlation, which precludes event ordering; simplistic time-out metrics for resetting the tracking of problems; ignoring individuals/sources responsible for exceptional activity) do not translate well to a malicious environment for detecting intrusions.

Earlier work in the intrusion-detection community attempting to address the issue of network surveillance includes the Network Security Monitor (NSM), developed at UC Davis [6], and the Network Anomaly Detection and Intrusion Reporter (NADIR) [7], developed at Los Alamos National Laboratory (LANL). Both performed broadcast LAN packet monitoring to analyze traffic patterns for known hostile or anomalous activity.[i] Further research by UC Davis in the Distributed Intrusion Detection System (DIDS) [23] and later Graph-based Intrusion Detection System (GRIDS) [25] projects has attempted to extend intrusion monitoring capabilities beyond LAN analysis, to provide multi-LAN and very large-scale network coverage.

This paper takes a pragmatic look at the issue of packet and/or datagram analysis based on statistical anomaly detection and signature-analysis techniques. This work is being performed in the context of SRI's latest intrusion-detection effort, EMERALD, a distributed scalable tool suite for tracking malicious activity through and across large networks [20]. EMERALD introduces a building-block approach to network surveillance, attack isolation, and automated response. The approach employs highly distributed, independently tunable, surveillance and response monitors that are deployable polymorphically at various abstract layers in a large network. These monitors demonstrate a streamlined intrusion-detection design that combines signature analysis with statistical profiling to provide localized real-time protection of the most widely used network services and components on the Internet.

Among the general types of analysis targets that EMERALD monitors are network gateways. We describe several analysis techniques that EMERALD implements, and discuss their use in analyzing malicious, faulty, and other exceptional network activity. EMERALD's surveillance modules will monitor entry points that separate external network traffic from an enterprise network and its constituent local domains.[ii] We present these surveillance techniques as complementary to the filtering mechanisms of a large enterprise network, and illustrate their utility in directly enhancing the security and stability of network operations.

We first consider the candidate event streams that pass through network entry points. Critical to the effective monitoring of operations is the careful selection and organization of these event streams such that an analysis based on a selected event stream will provide meaningful insight into the target activity. We identify effective analytical techniques for processing the event stream given specific analysis objectives. Sections 4 and 5 explore how both statistical anomaly detection and signature analysis can be applied to identify activity worthy of review and possible response. All such claims are supported by examples. More broadly, in Section 6 we discuss the correlation of analysis results produced by surveillance components deployed independently throughout the entry points of our protected intranet. We discuss how events of limited significance to a local surveillance monitor may be aggregated with results from other strategically deployed monitors to provide insight into more wide-scale problems or threats against the intranet. Section 7 discusses the issue of response.

3. Event Stream Selection

The success or failure of event analysis should be quantitatively measured for qualities such as accuracy and performance: both are assessable through testing. A more difficult but equally important metric to assess is completeness. With regard to network surveillance, inaccuracy is reflected in the number of legitimate transactions flagged as abnormal or malicious (false positives), incompleteness is reflected in the number of harmful transactions that escape detection (false negatives), and performance is measured by the rate at which transactions can be processed. All three measurements of success or failure directly depend on the quality of the event stream upon which the analysis is based. Here, we consider the objective of providing real-time surveillance of TCP/IP-based networks for malicious or exceptional network traffic. In particular, our network surveillance mechanisms can be integrated onto, or interconnected with, network gateways that filter traffic between a protected intranet and external networks.

IP traffic represents an interesting candidate event stream for analysis. Individually, packets represent parsable activity records, where key data within the header and data segment can be statistically analyzed and/or heuristically parsed for response-worthy activity. However, the sheer volume of potential packets dictates careful assessment of ways to optimally organize packets into streams for efficient parsing. Thorough filtering of events and event fields such that the target activity is concisely isolated, should be applied early in the processing stage to reduce resource utilization.

With respect to TCP/IP gateway traffic monitoring, we have investigated a variety of ways to categorize and isolate groups of packets from an arbitrary packet stream. Individual packet streams can be filtered based on different isolation criteria, such as

- *Discarded traffic:* packets not allowed through the gateway because they violate filtering rules.[iii]
- *Pass-through traffic:* packets allowed into the internal network from external sources.
- *Protocol-specific traffic:* packets pertaining to a common protocol as designated in the packet header. One example is the stream of all ICMP packets that reach the gateway.
- *Unassigned port traffic:* packets targeting ports to which the administrator has not assigned any network service and that also remain unblocked by the firewall.
- *Transport management messages:* packets involving transport-layer connection establishment, control, and termination (e.g., TCP SYN, RESET, ACK, [`window resize`]).
- *Source-address monitoring:* packets whose source addresses match well-known external sites (e.g., connections from satellite offices) or have raised suspicion from other monitoring efforts.
- *Destination-address monitoring:* all packets whose destination addresses match a given internal host or workstation.
- *Application-layer monitoring:* packets targeting a particular network service or application. This stream isolation may translate to parsing packet headers for IP/port matches (assuming an established binding between port and service) and rebuilding datagrams.

In the following sections we discuss how such traffic streams can be statistically and heuristically analyzed to provide insight into malicious and erroneous external traffic. Alternative sources of event data are also available from the report logs produced by the various gateways, firewalls, routers, and proxy-servers (e.g., router syslogs can in fact be used to collect packet information from several products). We explore how statistical and signature analysis techniques can be employed to monitor various elements within TCP/IP event streams that flow through network gateways. We present specific techniques for detecting external entities that attempt to subvert or bypass internal network services. Techniques are suggested for detecting attacks against the underlying network infrastructure, including attacks using corruption or forgery of legitimate traffic in an attempt to negatively affect routing services, application-layer services, or other network controls. We suggest how to extend our surveillance techniques to recognize network faults and other exceptional activity. We also discuss issues of distributed result correlation.

4. Traffic Analysis with Statistical Anomaly Detection

SRI has been involved in statistical anomaly-detection research for over a decade [1], [5], [10]. Our previous work focused on the profiling of user activity through audit-trail analysis. Within the EMERALD project, we are extending the underlying statistical algorithms to profile various aspects of network traffic in search of response- or alert-worthy anomalies.

The statistical subsystem tracks subject activity via one or more variables called *measures*. The statistical algorithms employ four classes of measures: categorical, continuous, intensity, and event distribution. *Categorical* measures are those that assume values from a categorical set, such as originating host identity, destination host, and port number. *Continuous* measures are those for which observed values are numeric or ordinal, such as number of bytes transferred. Derived measures also track the intensity of activity (that is, the rate of events per unit time) and the ``meta-distribution'' of the measures affected by recent events. These derived measure types are referred to as *intensity* and *event distribution*.

The system we have developed maintains and updates a description of a subject's behavior with respect to these measure types in a compact, efficiently updated *profile*. The profile is subdivided into short- and long-term elements. The short-term profile accumulates values between updates, and exponentially ages values for comparison to the long-term profile. As a consequence of the aging mechanism, the short-term profile characterizes the recent activity of the subject, where ``recent'' is determined by the dynamically configurable aging parameters used. At update time (typically, a time of low system activity), the update function folds the short-term values observed since the last update into the long-term profile, and the short-term profile is cleared. The long-term profile is itself slowly aged to adapt to changes in subject activity. Anomaly scoring compares related attributes in the short-term profile against the long-term profile. As all evaluations are done against empirical distributions, no assumptions of parametric distributions are made, and multi-modal and categorical distributions are accommodated. Furthermore, the algorithms we have developed require no *a priori* knowledge of intrusive or exceptional activity. A more detailed mathematical description of these algorithms is given in [9], [26].

Our earlier work considered the subject class of users of a computer system and the corresponding event stream the system audit trail generated by user activity. Within the EMERALD project, we generalize these concepts so that components and software such as network gateways, proxies, and network services can themselves be made subject classes. The generated event streams are obtained from log files, packet analysis, and--where required--special-purpose instrumentation made for services of interest (e.g., FTP, HTTP, or SMTP). As appropriate, an event stream may be analyzed as a single subject, or as multiple subjects, and the same network activity can be analyzed in several ways. For example, an event stream of dropped packets permits analyses that track the reason each packet was rejected. Under such a scenario, the firewall rejecting the packet is the subject, and the measures of interest are the reason the packet was dropped (a categorical measure), and the rate of dropped packets in the recent past (one or more intensity measures tuned to time intervals of seconds to minutes). Alternatively, these dropped packets may be parsed in finer detail, supporting other analyses where the subject is, for example, the identity of the originating host.

EMERALD can also choose to separately define satellite offices and ``rest of world'' as different subjects for the same event stream. That is, we expect distinctions from the satellite office's use of services and access to assets to deviate widely from sessions originating from external nonaffiliated sites. Through satellite session profiling, EMERALD can monitor traffic for signs of unusual activity. In the case of the FTP service, for example, each user who gives a login name is a subject, and ``anonymous'' is a subject as well. Another example of a subject is the network gateway itself, in which case there is only one subject. All subjects for the same event stream (that is, all subjects within a subject class) have the same measures defined in their profiles, but the internal profile values are different.

As we migrate our statistical algorithms that had previously focused on user audit trails with users as subjects, we generalize our ability to build more abstract profiles for varied types of activity captured within our generalized notion of an event stream. In the context of statistically analyzing TCP/IP traffic streams, profiling can be derived from a variety of traffic perspectives, including profiles of

- Protocol-specific transactions (e.g., all ICMP exchanges)
- Sessions between specific internal hosts and/or specific external sites
- Application-layer-specific sessions (e.g., anonymous FTP sessions profiled individually and/or collectively)

- Discarded traffic, measuring attributes such as volume and disposition of rejections
- Connection requests, errors, and unfiltered transmission rates and disposition

Event records are generated either as a result of activity or at periodic intervals. In our case, activity records are based on the content of IP packets or transport-layer datagrams. Our event filters also construct interval summary records, which contain accumulated network traffic statistics (at a minimum, number of packets and number of kilobytes transferred). These records are constructed at the end of each interval (e.g., once per N seconds).

EMERALD's statistical algorithm adjusts its short-term profile for the measure values observed on the event record. The distribution of recently observed values is evaluated against the long-term profile, and a distance between the two is obtained. The difference is compared to a historically adaptive, subject-specific deviation. The empirical distribution of this deviation is transformed to obtain a score for the event. Anomalous events are those whose scores exceed a historically adaptive, subject-specific score threshold based on the empirical score distribution. This nonparametric approach handles all measure types and makes no assumptions on the modality of the distribution for continuous measures.

The following sections provide example scenarios of exceptional network activity that can be measured by an EMERALD statistical engine deployed to network gateways.

4.1 Categorical Measures in Network Traffic

Categorical measures assume values from a discrete, nonordered set of possibilities. Examples of categorical measures include

- Source/destination address: One expects, for example, accesses from satellite offices to originate from a set of known host identities.
- Command issued: While any single command may not in itself be anomalous, some intrusion scenarios (such as ``doorknob rattling") give rise to an unusual mix of commands in the short-term profile.
- Protocol: As with commands, a single request of a given protocol may not be anomalous, but an unusual mix of protocol requests, reflected in the short-term profile, may indicate an intrusion.
- Errors and privilege violations: We track the return code from a command as a categorical measure; we expect the distribution to reflect only a small percent of abnormal returns (the actual rate is learned in the long-term profile). While some rate of errors is normal, a high number of exceptions in the recent past is abnormal. This is reflected both in unusual frequencies for abnormal categories, detected here, and unusual count of abnormal returns, tracked as a continuous measure as described in Section 4.2.
- Malformed service requests: Categorical measures can track the occurrence of various forms of bad requests or malformed packets directed to a specific network service.
- Malformed packet disposition: Packets are dropped by a packet filter for a variety of reasons, many of which are innocuous (for example, badly formed packet header). Unusual patterns of packet rejection or error messages could lead to insight into problems in neighboring systems or more serious attempts by external sites to probe internal assets.
- File handles: Certain subjects (for example, anonymous FTP users) are restricted as to which files they can access. Attempts to access other files or to write read-only files appear anomalous Such events are often detectable by signature analysis as well.

The statistical component builds empirical distributions of the category values encountered, even if the list of possible values is open-ended, and has mechanisms for ``aging out" categories whose long-term probabilities drop below a threshold.

The following is an example of categorical measures used in the surveillance of proxies for services such as SMTP or FTP. Consider a typical data-exchange sequence between an external client and an internal server within the protected network. Anonymous FTP is restricted to certain files and directories; the names of these are categories for measures pertaining to file/directory reads and (if permitted) writes. Attempted accesses to unusual directories appear anomalous. Monitors dedicated to ports include a categorical measure whose values are the protocol used. Invalid requests often lead to an access violation error; the type of error associated with a request is another example of a categorical measure, and the count or rate of errors in the recent past is tracked as continuous measures, as described in Section

4.2 Continuous Measures in Network Traffic

Continuous measures assume values from a continuous or ordinal set. Examples include inter-event time (difference in time

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.

WHAT WILL YOU BUILD? | sales@docketalarm.com | 1-866-77-FASTCASE

fastcase®
Smarter legal research.