

ADEPTS: Adaptive Intrusion Response using Attack Graphs in an E-Commerce Environment

Bingrui Foo, Yu-Sung Wu, Yu-Chun Mao, Saurabh Bagchi, Eugene Spafford⁺
Center for Education and Research in Information Assurance and Security (CERIAS)
Dependable Computing Systems Laboratory (+) School of Computer Sciences
School of Electrical and Computer Engineering Purdue University
Purdue University
E mail: {foob, yswu, maoy, sbagchi, spaf}@purdue.edu
Contact author: Saurabh Bagchi

Abstract

Distributed systems with multiple interacting services, such as distributed e-commerce systems, are suitable targets for malicious attacks because of the potential financial impact. Intrusion detection in such systems has been an active area of research, while the problem of automated response has received relatively less attention. The thought often is that a system administrator will be included in the loop for troubleshooting once the alert about a possible intrusion has been raised. In this paper, we present the design of automated response mechanisms in an intrusion tolerant system called ADEPTS. The particular type of response we focus on enforces containment in the system, through which it localizes the effect of the intrusion thus allowing the system to provide service, albeit degraded. Containment can be very important in a large class of distributed systems in which a single compromised service can affect other services through the mutual interactions. ADEPTS uses a graph of intrusion goals, called I-GRAPH as the underlying representation in the system. In response to alerts from an intrusion detection framework, ADEPTS executes an algorithm to determine the possible path of spread of the intrusion and the appropriate response to deploy. A feedback mechanism evaluates the success of a deployed response and uses that in guiding future choices. ADEPTS is demonstrated on a distributed e-commerce system and evaluated using the survivability metric whose value depends on the operational services in the face of an intrusion.

Keywords: automated intrusion response, intrusion containment, e-commerce system, survivability, dependency graphs.

1 Introduction

Distributed systems comprising multiple services interacting among themselves to provide end-user functions are becoming an increasingly important platform for business-to-business (B2B) and business-to-consumer (B2C) systems. As an example, electronic commerce, or e-commerce, has been touted as the next wave in the Internet

revolution. The huge financial stakes involved in e-commerce make the distributed system infrastructure supporting e-commerce prime candidates for computer security attacks.

Such motivations have long led to interest in securing distributed systems through detection of intrusions. This is typically achieved by analyzing the signatures of incoming packets and either matching them against known attack patterns (misuse based signatures), or against patterns of expected system behavior (anomaly based signatures). In order to meet the challenges of always-on, on-demand service availability, an e-commerce system needs to be resilient to security attacks. Resilience must include both intrusion detection and intrusion response. Compared to the problem of detection, automated response has received far less research attention. This has typically been considered the domain of system administrators who manually "patch" the system in response to detected attacks. However, as networked e-commerce services become ubiquitous and they are often placed in environments difficult to reach for human intervention, automated tools for intrusion response gain importance.

The rudimentary response mechanism often bundled with anti-virus or intrusion detection system (IDS) products [36],[37] overwhelmingly consider only immediate local responses that are directly suggested by the detected symptom. For example, a file being infected with a virus may cause the anti-virus product to quarantine the file and disable all access to the file [48],[49], or a suspect packet being flagged by a network IDS may cause the specific network connection to be terminated. While these may be applicable in stand-alone systems, they do not account for interaction effects among multiple components. The few available dedicated intrusion response systems are found to be lacking in one or more dimensions that make them unsuitable for protecting dynamic and complex distributed systems. Some of the commonly observed shortcomings are the system may have a static mapping of symptoms from the detector to the response, may not take feedback into account for determining future responses, may assume perfect detectors with no missed and no false alarms, or may assume perfect success rate for a deployed response. The complex interactions and the complex software running the distributed applications, the non-determinism in the execution environment, and the reality of new forms of intrusions surfacing would make any one of the above shortcomings fatal for an intrusion response system for a distributed enterprise.

In this paper, we focus on one of the most important kinds of automated response, namely, *containment*. Containment implies restricting the effect of the intrusion to a subset of the entire set of services, which may allow users access to limited functionality of the system. For example, browsing a store catalog and checking on a previously placed order may be available, while placing new orders may not be. There are several challenges to the problem of containment. First, the systems often have close coupling between the services with frequent interactions of different kinds, such as read, write, and execute. This allows a compromised service to spread the effect to multiple services. A second challenge is that the existing interactions between e-commerce system components should not be substantially altered during normal execution in order to enforce containment during periods of intrusion. Examples of unacceptable change may include mandating interactions pass through additional checks inlined in the execution path, intermediaries, or be executed over slower channels. Third, the system will have to consider the possibility of imperfect detectors providing false alarms or missing alarms, and imperfect response actions, which do not have 100% coverage.

In this paper, we present the design and implementation of an Adaptive Intrusion Tolerant System, *ADEPTS*, for containing intrusions in a distributed system of interacting services. *ADEPTS* uses an **Intrusion-Graph** (*I-GRAPH*) to represent paths for the intrusion to spread from one service to its neighbor. Alarms from a detection system, which may be off-the-shelf or from our previous work [6], are mapped to the *I-GRAPH* nodes. *ADEPTS* estimates the likely path of spread of the intrusion from the alarms and the structure of the *I-GRAPH* and then determines the appropriate response(s) to take. This decision is based on the disruptivity of the response to legitimate system activities, the previous success of the response, and the confidence that the determined intrusion is indeed taking place. The response has the goal of preventing the escalation of the intrusion and possible spread from one service to another. *ADEPTS* can function in multiple levels of “paranoia” depending on the policy level, from an aggressive mode with an elevated threat perception to a conservative mode.

The metric used to evaluate an intrusion tolerant system has to be carefully chosen. Low-level metrics, such as the latency of detection or false and missed alarm rates do not fully capture the effect of an intrusion on the system’s functionality. We propose the use of the metric called *survivability* [11] for evaluating the effect of *ADEPTS*. We define it such that its value depends on the set of high-level system transactions that can be achieved and the set of high-level system goals (e.g., keep users’ private information secure) that are *not* violated in the

event of an intrusion. A high level transaction relies on certain chains of interactions between multiple services being functioning. Preserving a high level goal implies thwarting certain intrusion goals from being reached.

The design of ADEPTS is realized in an implementation which provides intrusion response service to a realistic distributed e-commerce system. The e-commerce system mimics an online book store system and two auxiliary systems for the warehouse and the bank. Real attack scenarios are injected into the system and ADEPTS' responses are deployed, which bring out the latency of the response action and the adaptive nature of ADEPTS. The survivability of the system is compared with no response mechanism, local responses only, and with ADEPTS.

We believe this paper breaks new ground in the following ways:

1. ADEPTS is the first system, to the best of our knowledge, that provides a structured methodology for containing intrusions in a distributed system. It is also the first system to aggregate the factors of severity of a response, its effectiveness, and the possibility of escalation to determine the appropriate set of responses.
2. ADEPTS can handle multiple concurrent alerts, imperfect detectors, and escalation due to failed response actions. It can also deal with unanticipated alerts and unknown vulnerabilities in the system components. Each of these is of critical importance in an intrusion tolerance system applied to a real-world system.
3. ADEPTS is demonstrated on a realistic distributed testbed with realistic transactions and attack scenarios.

However, the work presented here *does not* have as its goal any of the following: intrusion detection for an e-commerce system, provide a methodology for structuring or composing an e-commerce system, design novel response actions for specific services in an e-commerce system, or provide a shrink-wrapped intrusion tolerance system to make an e-commerce system resilient to specific classes of attacks.

The rest of the paper is organized as follows. Section 2 refers to related research. Section 3 presents the design of ADEPTS. Section 4 describes the implementation and the e-commerce testbed on which ADEPTS is deployed. Section 5 presents the experiments and the results. Section 0 concludes the paper with mention of some future work.

2 Related Research

The devastating impact of computer security attacks to today's electronic world has spurred enormous interest in intrusion detection research, both from academic and commercial quarters. In order to guarantee the

requirement for continuous availability of the services, it is also important to consider how the system reacts once the intrusion is detected. The majority of current IDSs stops with flagging alarms and relies on manual response by the security administrator or system administrator. This results in delays between the detection of the intrusion and the response which may range from minutes to months. Cohen showed using simulated attack scenarios that given a ten hour delay from detection, 80% of the attacks succeed and given thirty hours, almost all the attacks succeed irrespective of the skill level of the defending system's administrator [8]. This insight has led to research in survivable systems engineering pioneered by CERT at CMU [9] and followed by several other researchers [26],[27],[28]. Survivability is loosely defined as the capability of a system to fulfill its mission, in a timely manner, in the presence of attacks, failures, or accidents ([10],[11]). The researchers identify the four key properties of survivable systems, namely, resistance to attacks, recognition of attacks and damage, recovery of essential and full services after attack, and adaptation and evolution to reduce effectiveness of future attacks. The part of the ADEPTS system presented in this paper is motivated by the requirement to provide the second and the fourth properties.

Intrusion response systems (IRS) can be considered to cover the last three properties and are therefore suitable for comparison with ADEPTS. A majority of the IRSs are *static* in nature in that they provide a set of preprogrammed responses that the administrator can choose from in initiating a response, e.g., [16],[35],[37],[38]. This may reduce the time gap between detection and response, but still leaves a potentially unbounded window of vulnerability. The holy grail is an IRS that can respond to an attack automatically. A handful of systems provide adaptive responses. In [17], the authors propose a network model that allows an IRS to evaluate the effect of a response on the network services. The system chooses in a *greedy* manner the response that minimizes the penalty. There are some studies which present taxonomy of offensive and defensive responses to aid in selection of coherent responses in an automated response system ([13],[18],[19]). Cooperating Security Managers (CSM) [12] is a distributed and host-based intrusion detection and response system. CSM proactively detects intrusions and reactively responds to them using the Fisch DC&A taxonomy [13]. It uses the suspicion level of the user as the only determining factor in the choice of response. A second system called EMERALD ([14],[15]) uses two factors in determining the response – the amount of evidence furnished to support the claim of intrusion and the severity of the response. None of the systems uses record of the past performance of the intrusion detection

system as measured by the incidence of false positives and false negatives. None keeps track of the success or failure of the deployed response nor provide a framework for easily incorporating these factors in the automated response determination. Another adaptive IRS is the Adaptive, Agent-based Intrusion Response System (AAIRS) ([20][21]). The work provides a good framework on which the IRS can be built. However, it does not provide any of the system-level techniques and algorithms that will be required for the AAIRS to work in practice. There is some previous work on protecting distributed systems against flooding based distributed denial of service (DDoS) attacks in an automated manner through rate limiting ([22][23],[24],[25]).

Fault trees have been used extensively in root cause analysis in fault tolerant systems (see [29] for pointers). They have also been used to a limited extent in secure system design ([30],[31],[32]). We use an attack graph representation with nodes as intermediate goals since the same intermediate goals show up in several attack paths. Graph theoretic approaches to modeling the temporal nature of security attributes is found in [33],[34]. The notion of privilege graphs introduced in [34] has some similarity to our I-GRAPH. However, they represent only attacks launched by escalating the privilege level of the attacker and the arcs are marked with weights representing the difficulty of the privilege escalation. The weights are dependent on several factors, such as the expertise and resources of the attacker, and therefore difficult to predict.

3 Design of ADEPTS

3.1 Overview

The goal of ADEPTS is to monitor and track intrusions as they occur in real-time and deploy various wide-ranging responses to contain and restrict the spread of attacks in the system. The system is subdivided into ADEPTS and the *payload* system, which includes the embedded detectors. The deployment of ADEPTS requires no modification to the payload and no access to its source code. The I-GRAPH models the paths an attacker can traverse to reach certain goals that adversely affect the payload. Our motive in designing ADEPTS is to proactively prevent an attacker from moving from one attack goal node to another, by responding appropriately at specific nodes. Here we give a high-level description of the process flow shown in Figure 1. As an alert comes into ADEPTS, it is mapped to nodes in the I-GRAPH followed by the execution of the response determination algorithm.

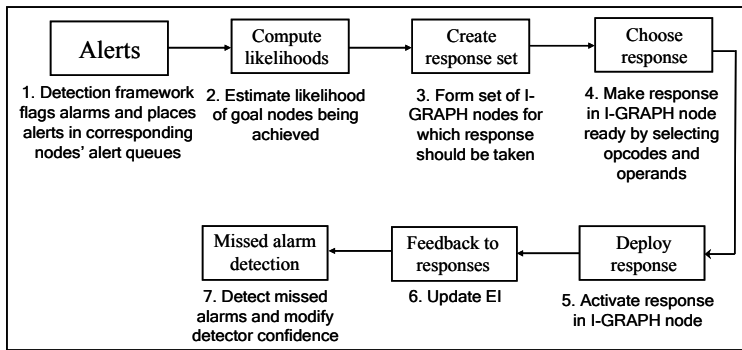


Figure 1: Overall process flow in ADEPTS showing the different stages starting from alert being flagged in the payload to response(s) being deployed

Throughout ADEPTS, three policy levels are used to control the behavior of the relevant algorithms – aggressive, moderate, and conservative. The three policies can be abstracted to represent a ratio of missed alarms to false alarms, with the aggressive policy having the lowest ratio and the conservative policy having the highest ratio.

3.2 I-GRAPH Structure

The I-GRAPH, is used as the underlying representation for knowledge about intrusions, as they spread achieving progressively wider set of goals.

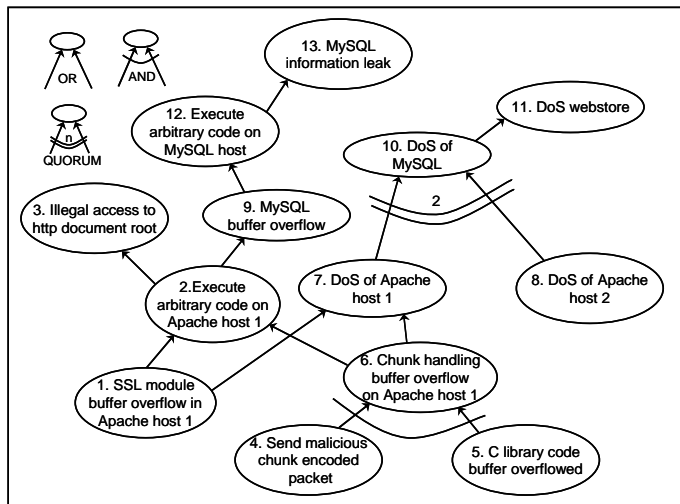


Figure 2: A section of the I-GRAPH referenced in later algorithms, with OR, AND, and Quorum edges present.

In the I-GRAPH representation, each intrusion goal is represented by one node in the graph. The final goal of the intrusion may be disrupting some high level system functionality, such as “Denial of service achieved against the online store”. This final step will be achieved through multiple small to moderate sized steps. A successful execution of a step is looked upon as achieving an intermediate intrusion goal and captured as an I-GRAPH node.

The intrusion goals have dependency relationships between one another. For example, in order to corrupt the data in the backend database server, one may need to exploit a vulnerability in the front-end web server. The edges are used to model this kind of dependency.

In the I-GRAPH, nodes are categorized into three types – OR node, AND node, and Quorum node. For an OR node to be achieved, any of its child nodes needs to be achieved, while for an AND node, all the child nodes have to be achieved. For a Quorum node, one can assign a Minimum Required Quorum (MRQ) on it, which represents the minimum number of child nodes whose goals need to be achieved in order for the Quorum node to be

achieved. Conforming to the traditional definition of quorums in fault tolerant systems, one may think MRQ as the minimum number of service replicas whose loss will affect the functionality of the service. An example fragment of the I-GRAPH used in our payload system, a distributed e-commerce system, is shown in Figure 2.

3.3 I-GRAPH Generation

A key issue in the usability of ADEPTS is the ease with which the I-GRAPH can be generated and updated as system configuration changes or new vulnerabilities are brought to light. We employ a semi-automated method called Portable I-GRAPH Generation (PIG) for this. PIG requires two inputs – vulnerability descriptions and system services description (SNet). Of the two inputs, the SNet is target system dependent. This is a directed graph, in which each node represents an individual service in the target system and an edge from node A to node B represents an *intrusion-centric channel*. An intrusion-centric channel means if A is compromised, then the intrusion can spread to B through the channel. An intrusion-centric channel may be of five kinds – (i) DOS channel: if the source service is subjected to a successful DoS attack, then the destination service can also be subjected to DoS; (ii) Network channel: there is a network data connection between A and B; (iii) Shared file channel; (iv) Shared memory channel; (v) Super channel: which combines the functionality of all of the above. The SNet is currently manually created for the target system, though in the future, some tool which can perform service discovery and interaction discovery (each an area of current research [50]) can perform this task automatically.

The second input to PIG is the target independent vulnerability descriptions. Information on the vulnerabilities can be obtained by querying the common vulnerability databases, such as CERT, Bugtraq, and CERIAs-VDB. For use in PIG, the vulnerability is specified through four fields – (i) *Name*: which is primarily useful for human reference. (ii) *Affected service*: which gives the service(s) in the SNet affected by the vulnerability; (iii) *Manifestation*: this is a Boolean expression in disjunctive normal form composed of five elementary manifestations, namely, leaking of information, execution of arbitrary code, incorrect behavior of service, DoS, and service termination. (iv) *Dependent vulnerability and services*: which denotes the dependence on other vulnerabilities and services that have to be compromised to exploit this vulnerability. The vulnerability definitions are analogous to the virus definitions used in anti-virus products. They can be developed either by the ADEPTS developer or by a third party. The basic idea behind the I-GRAPH generation algorithm is that when a vulnerability

description is read in, a corresponding node in the I-GRAPH is created, thus creating a one-to-one map. In the next step, the algorithm checks for nodes in the I-GRAPH that this newly created node can get connected to. For this step, it relies on information from both the SNet and the vulnerability descriptions to decide whether spread of the intrusion is possible from the newly created node to the other nodes and vice-versa.

3.4 Algorithms for Determining Response Locations

3.4.1 CCI Computation Algorithm

The goal of the algorithm is to determine, based on the received alerts from the detectors, which of the I-GRAPH goal nodes are likely to have been achieved. Each detector provides confidence values for its alerts, termed *alert confidence*. If the detector does not provide an inbuilt confidence value with the alert, then the alert confidence value is set to one. When a detector flags an intrusion, the alerts are placed in the I-GRAPH nodes with the corresponding intrusion event. The *Compromised Confidence Index* (CCI) of a node in the I-GRAPH is a measure of the likelihood that the node has been achieved. It is computed based on the alert confidence corresponding to the node and the CCI of its immediate children nodes. Mathematically, the CCI of a node is given by

$$CCI = \begin{cases} \text{alert confidence} & , \text{ nodes with no children} \\ f'(CCI_i) & , \text{ nodes with no detectors} \\ f(f'(CCI_i), \text{alert confidence}) & , \text{ otherwise} \end{cases} \quad f' = \begin{cases} \max(CCI_i) & , \text{ OR edges} \\ \min(CCI_i) & , \text{ AND edges} \\ \begin{cases} \text{Mean}(CCI_i \mid CCI_i > \tau_N) & , \text{ quorum met} \\ 0 & , \text{ quorum not met} \end{cases} \end{cases}$$

where CCI_i corresponds to the CCI of the i^{th} child.

where τ_N is a threshold per node.

The intuition is that for an OR edge, the parent node can be achieved if any of its children nodes is achieved and therefore the likelihood is the maximum of that of all of its children. For an AND edge, all the children nodes have to be achieved and therefore the likelihood is as much as the least likely child node. For Quorum edges, if the quorum is not met, then the higher goal is *not* achieved, but if met, the likelihood of it being achieved only depends on the children nodes that have achieved the quorum. The function f allows various weights to be assigned to determine the relative importance placed to the alert or the position of the node in the I-GRAPH. For simplicity, the function for the current design of ADEPTS is the statistical mean.

When new alerts arrive, the nodes corresponding to these alerts are passed to this algorithm, the I-GRAPH is traversed in breadth-first-search (BFS) order starting from the nodes corresponding to the new alerts, and the CCIs of the nodes are computed until each reachable node has been traversed at most once. The advantage of such a traversal allows us to convert the I-GRAPH into an I-DAG and thereby avoid cycling infinitely through the graph. The disadvantage of such a traversal is that some causal linkage of events between nodes will be ignored, which reduces the accuracy of the CCI computation. This disadvantage is not expected to be very significant since the temporal order of alerts is also often the causal order.

The alert confidence used to update the CCI is chosen based on policy. For an aggressive policy, the maximum alert confidence in the alert queue is used; for a moderate policy, the maximum of a subset of alert confidences based on the most recent alerts is chosen; for a conservative policy, the alert confidence corresponding to the most recent alert is chosen. The alert confidence provided by a detector has to be moderated by the confidence on the detector. ADEPTS has a mechanism to determine if a detector misses alarms and adjust the detector confidence accordingly. Qualitatively, if ADEPTS sees that for a given node n_i , its children nodes as well as parent nodes are flagged but n_i is not, then it anticipates probabilistically that the detectors have missed flagging the alert.

3.4.2 Response Set Computation Algorithm

The purpose of this algorithm is to determine the nodes in the I-GRAPH where current attacks will most likely spread to. This will allow the response algorithm to deploy appropriate responses at those locations. The I-GRAPH is traversed in reverse order of the CCI computation algorithm, continuing until all reachable nodes are traversed at most once. During the traversal, each node is labeled as one of: (i) *Strong Candidate* (SC), if $CCI > \tau_N$; (ii) *Weak Candidate* (WC), if $CCI \leq \tau_N$ but further traversal across only AND edges can reach a SC node; (iii) *Very Weak Candidate* (VWC), if $CCI \leq \tau_N$ but further traversal across any type of edge can reach a SC node; (iv) *Non-Candidate* (NC), otherwise. If the CCI of a node is computed to be greater than τ_N , the system concludes the node has been achieved. Therefore the SC label on a node is a strong indicator that the node has been achieved, while the WC or VWC label indicates smaller likelihoods of nodes being achieved due to evidence from their parents.

Next, some nodes are placed in a *response set*, indicating to the response system where responses should be deployed. For an aggressive policy, all SC nodes, and WC and VWC nodes which have at least one immediate

NC parent node are placed in the response set. For a moderate policy, all SC and WC nodes that have at least one immediate NC parent node are placed in the response set. For a conservative policy, all SC nodes that have at least one immediate NC parent node are placed in the response set. The aggressive, moderate, and conservative policies provide increasingly less disruption as well as protection.

Referencing Figure 2, suppose the CCIs of nodes 1, 4, 5, 6, and 7 are 0.8, 0.3, 0.4, 0.4, and 0.7, respectively, and the other nodes have negligible CCI. Then a possible partial traversal order would be 7, 6, 1, 5, and 4. Suppose the threshold $\tau_N = 0.5$, then node 7 will be a SC node, nodes 6, 1, 5, 4 VWC nodes, and all other nodes NC nodes. For an aggressive policy, nodes 7 and 2 would be in the response set; for a moderate or conservative policy, only node 7 would be in the response set.

3.5 Response Process: Response Repository

The deployment of the response is achieved by a *Response Repository*, a *Response Control Center*, and distributed *Response Execution Agents*. The Response Repository stores the responses available for deployment in a payload system. Each response in the repository consists of an opcode and one or more operands, with wildcards allowed for each. The opcode is the response command, and the operands are the different parameters that need to be specified in order to execute the response. For example, the opcode for the response command of dropping incoming packets from a remote IP to a local port is DROP_INPUT, and the corresponding operands are REMOTE_IP and LOCAL_PORT. The opcode and the operands together make up a complete response command. The response structure allows ADEPTS fine-grained customization of the available responses.

3.6 Response Control Center

The opcode is selected based on the ability of the opcode to cut off the *attack-centric channels* as defined in Section 3.4. The Response Set Computation algorithm (Section 3.4.2) sends to the Response Control Center the list of I-GRAPH nodes, which are candidates for the deployment of responses. For each node, the Response Control Center selects a set of candidate response opcodes that can be used to prevent attacks from spreading via the node's outgoing intrusion-centric channels. The choice is determined by the type of the channel. For example, the file access based opcodes, such as DENY_FILE_ACCESS or DISABLE_WRITE, are selected as candidate response opcodes if an outgoing shared file channel is present.

After the opcodes have been chosen, the Response Control Center generates a list of complete response commands by collecting suitable operands. For this, it examines the alert events stored in the *alert queue* of the I-GRAPH node and uses them to fill in the operands that are required by the selected opcodes. An opcode can be combined with multiple operands during this phase. For example, for an opcode KILL_PROCESS, the control center may extract PID#1 from alert event#1 and PID#2 from alert event #2, both in the alert queue. Then, the response command KILL_PROCESS PID#1, PID#2 is generated for subsequent evaluation.

3.6.1 *Pick Responses to Deploy*

For each selected response command, the Response Control Center computes the Response Index (RI). The RI takes into the account the estimated effectiveness of the response to the particular attack, measured by the Effectiveness Index (EI), and the perceived disruptiveness of the response to legitimate users of the system, measured by the Disruptiveness Index (DI). The EI and the DI are both specific to the response command (opcode-operand combination) and the node in the I-GRAPH to which the response is mapped. The RI is given by $RI = a.EI - b.DI$, where a and b are deployment parameters.

Note that EI of an identical response command may differ for different attacks that map to different I-GRAPH nodes. For example, blocking port 65000 or 16660 may be useful to block the *stacheldraht* DDoS attack tool [42] but is unlikely to be effective for the TFN DDoS attack tool [43]. The control center chooses the response with the highest RI among the candidate responses, with a threshold being used to suppress a response that falls below it. This ensures that ineffective or highly disruptive responses are not deployed. If no response is chosen for a particular node, then the next higher level node is searched for possible responses.

In the event that the payload system is under multiple concurrent attacks, multiple alert events are mapped to the same I-GRAPH node. ADEPTS deploys responses for different alerts received in a short span of time, which may correspond to each individual attack instance. A heuristic to distinguish different concurrent attacks, involving clustering source IP addresses, destination IP addresses, source ports, destination ports, user accounts and initiated processes, proposed in [21], can be easily integrated into ADEPTS.

3.6.2 *Contradiction, Equivalence, Subset, and Super Set Relation between Responses*

Before initiating execution of the chosen responses, the Response Control Center identifies the relations between the active responses and the pending responses and the validity of the response itself. The newly selected

response is suppressed and a new response command with the second highest RI is considered if one of the following conditions holds: the new command is a subset or the equivalent of an active response, the new response contradicts an active response, or an *inconsistent* response has been generated. An example of an inconsistent response is the command to block incoming UDP packets toward port 80 is inconsistent since HTTP packets directed to port 80 transmit under TCP protocol. If there is overlap between the new response and an active response, the ideal strategy would be to deploy the non-overlapping part of the new response. Since it is difficult to extract subsets out of a response in an automated manner, we enforce the design choice on the responses in ADEPTS that they are all non-overlapping. This is possible to achieve because of the fine granularity of the responses.

3.6.3 Handling unknown alerts

In a real-world deployment, it is quite probable that the I-GRAPH for the payload system is incomplete. Thus, ADEPTS would be unable to map an incoming alert from a detector to a node in the I-GRAPH. To handle this situation, ADEPTS has the provision of a general I-GRAPH node per host. The alert would be mapped to the general I-GRAPH node for the host that is the destination of the attack as is easily deducible from the alert. Since the general node represents unknown vulnerabilities, it is not connected to any other nodes in the I-GRAPH. Therefore, ADEPTS cannot follow the strategy of traversing I-GRAPH edges as in Section 3.6.1. In this case, the Response Control Center can simply report the instance to the administrator and take one of a set of pre-specified *general responses*. The *general responses* are the commands that would be possible to deploy with very little knowledge of the operands, such as killing a process (need process ID), shutting down a service (need service ID), or restarting a host (need host ID).

3.7 Deploying Response & Providing Feedback

Feedback to the response system is crucial for ADEPTS, providing the runtime mechanism to bias response choices in favor of those that have been effective in the past. This feedback is provided by dynamically varying the EI of the response. After a response has been deployed by ADEPTS, the feedback system checks to see if any active response action is deployed on an edge that can be used to reach a node in the currently computed response set. If such a response action exists, it is indication that the response action possibly failed and its EI is decreased.

The amount by which the EI of the response is decreased depends on whether the response is on an AND edge, OR edge, or Quorum edge to the node in the response set. If it is on an AND edge, then it is certain that the response failed and thus the node was achieved. Therefore, the EI is decreased by a fixed fraction for responses on all the edges. If the response is on an OR or Quorum edge, then the EI is decreased in the proportion of the CCI values of the nodes, the maximum decrease being the same as in the AND case. When a response's Time To Live (TTL) expires or when an administrator manually deactivates a response, the EI of the response action is increased by a fixed percentage under the intuition that the response was successful since further alerts were not observed.

Referencing Figure 2, suppose an active response is present on the edge between node 1 and 7, and node 10 is in the response set. Also suppose the CCI of nodes 1, 6, 7, and 8 are 0.8, 0.3, 0.6, 0.1 respectively and the fixed fraction to decrease is 0.2. Then for the active response, $EI_{new} = EI_{old} - 0.2 \times \frac{0.6}{0.6+0.1} \times \frac{0.8}{0.8+0.3} = EI_{old} - 0.125$.

4 Implementation of ADEPTS & Testbed

4.1 Description of E-Commerce Application

Figure 3 depicts the testbed that we use for experiments on ADEPTS. The payload system mimics an e-Commerce webstore, which has two Apache web servers running webstore applications, which are based on Cubecart [44] and are written in the PHP scripting language. In the backend, there's a MySQL database which stores all the store's information, which includes products inventory, products description, customer accounts, and order history. There are two other organizations with which the webstore interacts – a Bank and a Warehouse. The Bank is a home-grown application which verifies credit card requests from the webstore. The Warehouse is also a home-grown application, which takes shipping requests from the webstore, checks inventory, applies charges on the customer's credit card account, and ships the product. The clients submit transactions to the webstore through a browser. Some important transactions are given in Table 1.

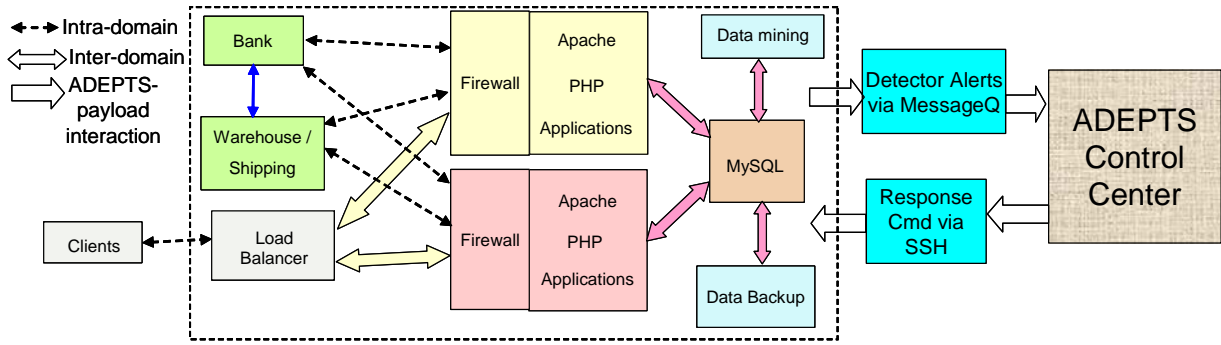


Figure 3: Layout of e-commerce testbed for the experiments on ADEPTS

Name	Description	Services involved	Weight
Browse webstore	Customer uses web browser to access webstore and browse the products available	Apache, MySQL	10
Add merchandise to shopping cart	Customer adds products to shopping cart	Apache, MySQL	10
Place order	Customer can input credit card information, submit orders, and webstore will authenticate credit card with bank	Apache, MySQL, bank	10
Charge credit card	Warehouse charges credit card through bank when order is shipped	Warehouse, bank	5
Admin work	Admins/webmasters can modify various source codes	Variable	10

Table 1: List of important transactions in e-commerce system

We set certain security goals for the system, the complement of which are specified in Table 2, along with the weights. Thus adding the word “prevent” before each gives the goal. The attached weights to the transactions and security goals are used for survivability computation in Section 5.

Illegal read of file (20)	Corruption of MySQL database (70)	Unauthorized credit card charges (80)
Illegal write to file (30)	Confidentiality leak of customer information stored in MySQL database (100)	Cracked administrator password (90)
Illegal process being run (50)	Unauthorized orders created or shipped (80)	

Table 2: List of security goals for e-commerce testbed

4.2 Detectors

For our testbed, multiple detectors which communicate with ADEPTS through secure channels are used. We use two off-the-shelf detectors – Snort [35] and Libsafe [45], and create three home-grown detectors. Snort is used for detecting intrusion patterns in network traffic while Libsafe is used to detect buffer overflows in protected C-library calls. We create a kernel-based *File Access Monitor*, which can detect file access attempts of monitored processes and compare these access attempts against preset rules to detect illegitimate activity. Also, we create a *Transaction Response Monitor*, which monitors the transaction response time of the webstore using requests from the Apache Benchmark [46]. Finally, there is an *Abnormal Account Activity Detector* at the Bank, which detects abnormal account activities such as excessive number of credit card transactions on one account.

The detectors used are all imperfect ones, with the possibility of missed alarms and false alarms. The detectors are not customized to the attack scenarios that the system is tested with. For the off-the-shelf detectors, the rules are taken from the public distribution, while for the others, the rules are created by a researcher separate from the group that generates the attack scenarios.

4.3 Attack Scenarios

The ADEPTS deployment is tested with different attack scenarios classified into three categories – leaking information, illegal transaction, and DoS. Each attack scenario consists of a set of attack steps, with an ultimate high-level goal. Each step of the attack scenario may be detected by none, one, or more of the detectors. A detector vector with the elements (Snort, Libsafe, FileAccess, Bank Monitor., Transaction Response Monitor) is assigned to each step of the attack scenario. A ‘1’ means that step can be detected by the corresponding detector. We show in Table 3 one sample scenario from each category – Scenario 0 is leaking of user information in the database (Leaking information), Scenario 1 is placing unauthorized orders (Illegal transaction), and Scenario 8 is vandalize webstore (DoS).

Steps	Scenario 0	Scenario 1	Scenario 8
0	Exploit Apache mod buffer overflow. (1, 0, 0, 0, 0)	Use php_mime_split (CVE-2002-0081) buffer overflow to insert malicious code into Apache. (1, 0, 0, 0, 0)	ModSSL Buffer overflow in Apache. (1, 0, 0, 0, 0)
1	Insert malicious code. (0, 0, 0, 0, 0)	'ls' to list webstore document root and identify the script code informing the warehouse to do shipments. (0, 0, 1, 0, 0)	A shell is created with Apache privilege. (0, 0, 1, 0, 0)
2	Ip/port scanning to find vulnerable SQL server. (1, 0, 0, 0, 0)	Send shipping request to warehouse and craft the request form so that a warehouse side buffer overrun bug fills the form with a victim's credit card number. (0, 1, 0, 0, 0)	Issue crontab command to exploit a vulnerability in cron daemon for creating a root privilege shell. (0, 0, 1, 0, 0)
3	Buffer overflow MYSQL to create a shell (/bin/sh). (0, 0, 1, 0, 0)	Unauthorized orders are made. (0, 0, 0, 1, 0)	Root privilege shell created out of the vulnerable cron daemon. (0, 0, 0, 0, 0)
4	Use malicious shell to steal information stored in MySQL. (0, 0, 1, 0, 0)		Corrupt the data stored in web server document root. (0, 0, 1, 0, 0)

Table 3: Attack steps for three attack scenarios used in experiments with ADEPTS

We also test ADEPTS with other attack scenarios involving buffer overflow attacks to steal client info, and other DoS attack scenarios entailing memory exhaustion in the Apache mime handling components or DDoS through huge number of legitimate transactions, such as product search. The entire I-GRAPH automatically generated by the PIG algorithm consists of 57 nodes and 1148 edges and is too large to be shown. A fragment of the I-GRAPH has been shown in Figure 2.

4.4 Response Repository for E-Commerce Testbed

Four types of response commands are included in the Response Repository – *general*, *file*, *network*, and *denial-of-service* types. The *general-type* commands can be deployed to block any types of *intrusion-centric* channels in the I-GRAPH, corresponding to the super channel. The other types of commands have a one-to-one map to the kinds of intrusion channels introduced in Section 3.3. The implementation of the file-type commands is achieved by using the Linux Intrusion Detection System (LIDS) 2.2.0 [38]. The implementation of the network-type commands is performed by using *iptables* [47]. The general type commands are killing a process and restarting or shutting down a service or a host. The file-type commands are to deny any access to a file, or selectively disable read, write, or execute access. The network-type commands are to block incoming or outgoing network connections, parameterized by source or destination port, IP, or protocol. The DOS-type commands are to limit the rates of various types of packets, such as SYN, ICMP echo, ICMP host not reachable, and SYN-ACK.

5 Experiments and Results

We perform three sets of experiments demonstrating the following (i) effect of attack scenarios on survivability with and without ADEPTS, (ii) the ability of ADEPTS to deploy responses as the speed of propagation of the attack varies, (iii) the adaptation in ADEPTS in choosing responses. All these experiments are conducted with ADEPTS using moderate policy with actual attack scenarios on the e-commerce testbed. For experiment 1 and 2, we define the survivability based on the high level transactions and security goals. The metric thus shows the effect of ADEPTS on the high level functioning of the e-commerce system.

$$\text{Survivability} = 1000 - \sum \text{unavailable transactions} - \sum \text{failed security goals} .$$

When a transaction became unavailable or the security goal is violated, the survivability drops by its corresponding weight, which was given in Table 1 and Table 2. Transactions become unavailable due to ADEPTS responses, such as rebooting a host, or attacks. Security goals may be violated due to the successful execution of an attack step. If a security goal is violated multiple times during an attack, then each violation causes a decrease in the survivability.

5.1 Experiment 1: Effect of Attack Scenarios on Survivability

The goal of experiment 1 is to show the comparative performance of ADEPTS in maintaining the survivability of the e-commerce system with respect to having no responses and only local responses. Three different attack scenarios are executed and the survivability calculated at each step of the attack scenario. For the local response case, the responses that came with the deployed detectors are used – Snort (IP blocking) and bank monitor (freeze credit card).




Figure 4: Survivability with attack steps with ADEPTS, local, and no response, and no response (a) scenario 0; (b) scenario 1

For the *leak of information* attack (Figure 4(a)), ADEPTS far outperforms the other two. The File Access Monitor detects a malicious shell being created with Apache privileges while Snort detects an Apache SSL module buffer overflow packet. Consequently, ADEPTS deploys aggressive responses to kill the process and block all following incoming packets from the attacker. The inability of the local response implemented by Snort to drop the IP packets in time causes the attack to continue to spread. For the *illegal transaction* attack (Figure 4(b)), the performance of the local response is noticeably worse than ADEPTS. ADEPTS deploys a successful response of disallowing shell commands with Apache privileges, earlier than the local response at the bank monitor. For the *distributed denial of service* attack (Figure 5), the graph shows the inability of any of the setups to respond effectively to the attack. The responses deployed by ADEPTS to limit the overall incoming packet rate, such as, blocking packets from the IP address with the highest rate of packet transmission, allowed for a slight decrease in the effectiveness of the DoS.

5.2 Experiment 2: Effect of Propagation Speed on Survivability

The experiment inspects the relationship between the ability of ADEPTS to protect the payload system and the attack propagation speed. We vary the delay between the attack steps in scenario 0 between 0-7 s. This could simulate a variety of factors, such as the attacker's skill level, condition of the network, difficulty of an attack step, etc. Figure 6 shows the survivability with different attack propagation speeds. The legend '0004' means there is a delay of zero seconds before attack step 0 is begun, between steps 0 and 1, and 1 and 2, while there is a delay of 4 s for all subsequent steps. We see that ADEPTS performs well with delays ≥ 4 s since the attack is stopped in the very first step. If however, the first step has zero delay, then ADEPTS is only able to block the attack at a later step, leading to a decreased survivability. With no delay at all between the steps, ADEPTS is only able to block the attack at the last step, which is still better than the no response case. In all cases, ADEPTS can maintain the survivability at a constant level once the blocking is successful.




Figure 5: Survivability with attack steps with ADEPTS, local response, and no response (scenario 4)




Figure 6: Effect of attack propagation speed on survivability with ADEPTS

5.3 Experiment 3: Adaptation of Response Mechanism in ADEPTS

Here we demonstrate the adaptive nature of ADEPTS through which it can change the response strategy as an attack escalates. Scenario 8, having 5 steps, with delay characteristic '00015' is used for demonstrating the process. This experiment is composed of multiple runs of scenario 8 (we show runs 1, 3, and 6). We see in Run 1, the initial choice of responses is poor ('X' by response implies failed) and the final step of scenario 8 is reached, where the attacker can create a root privilege shell and tamper with Apache's HTTP documents. In Run 3, after a series of EI tuning from Runs 1 and 2, the response of rebooting Apache is deployed in step 2. In Run 6, the response of rebooting Apache is further moved ahead to step 1. The farther ahead of the final attack goal (step 8)

that the attack is blocked, the higher is the safety that the goal has not been reached. Thus, the choice of responses in run 6 is the best and ADEPTS is seen to improve its response choices through observing multiple attacks.

R1 Block port 443 from attacker's src IP	R3 Block attacker's source IP	R5 Restart Aapache	R7 Deny access to crontab command and kill process if still running
R2 Kill the Apache privilege shell	R4 Kill crontab process	R6 Reboot Apache's host machine	R8 Set Apache's HTTP document directory to READONLY

Run 1	
Step	Attack impacts
0	Apache MOD_SSL Buffer Overflow
1	Apache Privilege Shell Created
	X R1 X R2
2	Executing crontab command
EI [R1] 1.1 → 1.072497 EI [R2] 1.1 → 1.058745	X R3 X R4
3	Put malicious data into Apache user's crontab
EI [R2] 1.058745 → 1.024404 EI [R4] 1.1 → 1.064321 EI [R3] 1.1 → 1.076214 EI [R1] 1.072497 → 1.049305	X R5
4	Root privilege Shell created out of cron daemon
5	Via the root shell, the attacker tampers with files under http document directory.
EI [R2] 1.024404 → 0.99328 EI [R4] 1.064321 → 1.031984 EI [R5] 1.1 → 1.077720 EI [R3] 1.076214 → 1.054415 EI [R1] 1.049305 → 1.028052	O R6 X R7 X R8
Attack Stopped	
EI [R8] 1.1 → 1.21 EI [R7] 1.1 → 1.21 EI [R6] from 1.1 → 1.21	

Run 3	
Step	Attack impacts
0	Apache MOD_SSL Buffer Overflow
1	Apache Privilege Shell Created
	X R1 X R7
2	Executing crontab command
EI [R7] 1.21 → 1.164620 EI [R1] 0.96081 → 0.936787	O R6 X R4
3	Put malicious data into Apache user's crontab
EI [R7] 1.16462 → 1.126844 EI [R4] 1.1 → 1.064321 EI [R1] 0.936787 → 0.916530 EI [R6] 1.331 → 1.302219	X R3
Attack Stopped	
EI [R3] up from 1.01072 → 1.111792	
Run 6	
0	Apache MOD_SSL Buffer overflow
1	Apache privilege shell created
	O R6 O R7
Attack Stopped	
EI [R7] 1.1 → 1.21 EI [R6] 1.401466 → 1.541612	

6 Conclusions

In the paper we have presented the design and implementation of an automated intrusion containment system called ADEPTS. ADEPTS uses a graph of intrusion goals called I-GRAPH. It provides a method to determine the possible path of spread of the intrusion, appropriate services where to deploy the response, and appropriately choose the response. ADEPTS is demonstrated on an e-commerce system with real attack scenarios.

We are currently investigating ways to synthesize new responses at runtime from the repository, designing protocols for handling concurrent attacks, distinguishing between attacks, and evaluating the convergence of its adaptation feature.

References

- [1] S. Garfinkel and G. Spafford, "Web Security & Commerce," O'Reilly, 1997.
- [2] A. Ghosh, "E-Commerce Security," John Wiley and Sons, Inc., Third Avenue, New York, 1998.
- [3] V. Hassler, "Security Fundamentals for E-commerce," Artech House, 2001.
- [4] L. D. Stein, "Web Security: A step-by-step reference guide," Addison Wesley, Reading, Massachusetts, 1999.
- [5] Forrester Research Inc., "US E-Commerce Overview: 2003 to 2008," Techstrategy Brief Series, July 2003.
- [6] Yu-Sung Wu, Bingrui Foo, Yongguo Mei, and Saurabh Bagchi, "Collaborative Intrusion Detection System (CIDS): A Framework for Accurate and Efficient IDS," In Proceedings of the 19th Annual Computer Security Applications Conference (ACSAC '03), December 8 - 12, 2003.
- [7] AT&T Research, "Graphviz - open source graph drawing software," Available at: <http://www.research.att.com/sw/tools/graphviz/>
- [8] F. B. Cohen, "Simulating Cyber Attacks, Defenses, and Consequences," Available at <http://all.net/journal/ntb/simulate/simulate.html>, May 13, 1999.
- [9] Carnegie Mellon, Software Engineering Institute, "Survivable Network Technology," Available at: <http://www.sei.cmu.edu/organization/programs/nss/surv-net-tech.html>
- [10] R. H. Anderson, A. C. Hearn, and R. O. Hundley, "Studies of Cyberspace Security Issues and the Concept of a U.S. Minimum Essential Information Infrastructure," Proceedings of the 1997 Information Survivability Workshop, CERT Coordination Center, Software Engineering Institute, Carnegie Mellon University, 1997.
- [11] R. Ellison, R. Linger, T. Longstaff, N. Mead, "Case Study in Survivable Network System Analysis," CMU/SEI-98-TR-014, ADA355070, Software Engineering Institute, Carnegie Mellon University, 1998.
- [12] G. B. White, E. A. Fisch, and U. W. Pooch, "Cooperating Security Managers: A Peer-based Intrusion Detection System," IEEE Network, vol. 10, no. 1, January/February, 1996, pp. 20-23.
- [13] E. A. Fisch, "Intrusion Damage Control and Assessment: A Taxonomy and Implementation of Automated Responses to Intrusive Behavior," Ph.D. Dissertation, Texas A&M University, College Station, TX, 1996.
- [14] P. A. Porras and P. G. Neumann, "EMERALD: Event Monitoring Enabling Responses to Anomalous Live Disturbances," in Proc. 20th National Information Systems Security Conf., Baltimore, MD, October 7-10, 1997, pp. 353-365.
- [15] P. G. Neumann and P. A. Porras, "Experience with EMERALD to Date," in Proc. 1st USENIX Workshop on Intrusion Detection and Network Monitoring, Santa Clara, CA, April 11-12, 1999.
- [16] T. Bowen, D. Chee, M. Segal, R. Sekar, T. Shanbhag, P. Uppuluri, "Building Survivable Systems: An Integrated Approach based on Intrusion Detection and Damage Containment," DARPA Information Survivability Conference and Exposition (DISCEX '00), pp. 84-99, vol. 2, Jan 2000.
- [17] Thomas Toth and Christopher Kruegel, "Evaluating the Impact of Automated Intrusion Response Mechanisms," 18th Annual Computer Security Applications Conference (ACSAC '02), December 9 - 13, 2002.
- [18] Curtis A. Carver, Jr., and Udo W. Pooch, "An Intrusion Response Taxonomy and its Role in Automatic Intrusion Response," Proceedings of the 2000 IEEE Workshop on Information Assurance and Security, United States Military Academy, West Point, NY, 6-7 June, 2000.
- [19] U. Lindqvist and E. Jonsson, "How to Systematically Classify Computer Security Intrusions," Proc. 1997 IEEE Symposium on Security and Privacy, Oakland, CA, May 4-7, 1997, pp. 154 - 163.
- [20] Daniel Ragsdale, Curtis Carver, Jeffery Humphries, and Udo Pooch, "Adaptation Techniques for Intrusion Detection and Intrusion Response Systems," In Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics, Nashville, Tennessee, October 8-11, 2000, pp. 2344-2349.
- [21] Curtis A. Carver, John M.D. Hill, and Udo W. Pooch, "Limiting Uncertainty in Intrusion Response," Proceedings of the 2001 IEEE Workshop on Information Assurance and Security United States Military Academy, West Point, NY, 5-6 June, 2001.
- [22] Recourse Technologies, "ManHunt product description," Available at: <http://www.recourse.com/products/manhunt/features.html>.
- [23] Dan Schnackenberg, Harley Holliday, Randall Smith, Kelly Djahandari, Dan Sterne, "Cooperative Intrusion Traceback and Response Architecture (CITRA)," DARPA Information Survivability Conference and Exposition (DISCEX II'01), June 12 - 14, 2001.
- [24] D. Sterne, K. Djahandari, B. Wilson, B. Babson, D. Schnackenberg, H. Holliday, and T. Reid, "Autonomic Response to Distributed Denial of Service Attacks," In Proceedings of the 4th International Symposium on Rapid Advances in Intrusion Detection, RAID 2001, Davis, CA, USA, October 2001.
- [25] R. Mahajan, S. Bellovin, S. Floyd, J. Ioannidis, V. Paxson, and S. Shenker, "Controlling High Bandwidth Aggregates in the Network," AT&T Center for Internet Research at ICSI (ACIRI), DRAFT, February 5, 2001. Available at: <http://www.research.att.com/~smb/papers/DDOS-lacc.pdf>

- [26] K. Kyamakya, K. Jobman, M. Meincke, "Security and Survivability of Distributed Systems: An Overview," At the 21st Century Military Communications (MILCOM '00), 2000.
- [27] S. Jha, J. Wing, R. Linger, T. Longstaff, "Survivability Analysis of Network Specifications," In Workshop on Dependability Despite Malicious Faults, International Conference on Dependable Systems and Networks (DSN), June 2000.
- [28] Matti A. Hiltunen, Richard D. Schlichting, Carlos A. Ugarte, and Gary T. Wong, "Survivability through Customization and Adaptability: The Cactus Approach," DARPA Information Survivability Conference and Exposition (DISCEX 2000), pp. 294--307, January 2000.
- [29] Daniel P. Siewiorek and Robert S. Swarz, "Reliable Computer Systems – Design and Evaluation," Chapter 5: Evaluation Criteria, pg. 350, A. K. Peters Ltd., Third Edition.
- [30] P.J. Brooke and R.F. Paige, "Fault Trees for Security System Analysis and Design," Journal of Computers and Security, 22(3):256-264, Elsevier, May 2003.
- [31] Guy Helmer, Johnny Wong, Mark Slagell, Vasant Honavar, Les Miller, and Robyn Lutz, "A Software Fault Tree Approach to Requirements Analysis of an Intrusion Detection System," In Proceedings of the First Symposium on Requirements Engineering for Information Security, 2001.
- [32] B. Schneier, "Attack Trees: Modeling Security Threats," Dr. Dobbs Journal, December 1999, Available at: <http://www.counterpane.com/attacktrees-ddj-ft.html>.
- [33] M. Dacier, Y. Deswarte, and M. Kaaniche, "Quantitative Assessment of Operational Security: Models and Tools," LAAS Research Report 96493, May 1996 (Extended version of "Models and Tools for Quantitative Assessment of Operational Security," Proc. IFIP/SEC'96).
- [34] R. Ortalo, Y. Deswarte, M. Kaaniche, "Experimenting with Quantitative Evaluation Tools for Monitoring Operational Security," IEEE Transactions on Software Engineering, Volume: 25 , Issue: 5 , pp. 633-650, Sept.-Oct. 1999.
- [35] "Snort: open source network intrusion detection system," At: <http://www.snort.org>
- [36] "Snort Flexible Response Add-On," Available at: http://cerberus.sourceforge.com/~jeff/archives/snort/sp_respond2/
- [37] "Snort In-line," Available at: <http://sourceforge.net/projects/snort-inline/>
- [38] "The Linux Intrusion Detection System (LIDS) Project," Available at: <http://www.lids.org>
- [39] US-CERT, "US-CERT Vulnerability Notes Database," At: <http://www.kb.cert.org/vuls>
- [40] Mitre Corporation, "Common Vulnerabilities and Exposures," At: <http://www.cve.mitre.org/>
- [41] Transaction Processing Performance Council (TPC), "TPC-W: A transactional web e-commerce benchmark," At: <http://www.tpc.org/tpcw/>
- [42] "The Stacheldraht Distributed Denial of Service Attack Tool," At: <http://staff.washington.edu/dittrich/misc/stacheldraht.analysis.txt>
- [43] "The Tribe Flood Network Distributed Denial of Service Attack Tool," At: <http://staff.washington.edu/dittrich/misc/tfn.analysis.txt>
- [44] "CubeCart: eCommerce modified," At: <http://www.cubecart.com>
- [45] "Libsafe: Protecting Critical Elements of Stacks," At: <http://www.research.avayalabs.com/project/libsafe/>
- [46] "Apache HTTP server benchmarking tool," At: <http://httpd.apache.org/docs-2.0/programs/ab.html>
- [47] "Netfilter/Iptables project home page," At: <http://www.netfilter.org/>
- [48] "Norton Antivirus" Available at <http://www.symantec.com/nav/>
- [49] "McAfee VirusScan" Available at <http://www.mcafee.com>
- [50] A. Brown, G. Kar, A. Keller, "An Active Approach to Characterizing Dynamic Dependencies for Problem Determination in a Distributed Application Environment," IEEE/IFIP International Symposium on Integrated Network Management, pp. 377-390, 2001.