Shield: Vulnerability-Driven Network Filters for Preventing Known Vulnerability Exploits

Helen J. Wang, Chuanxiong Guo, Daniel R. Simon, and Alf Zugenmaier
{helenw, t-chuguo, dansimon, alfz}@microsoft.com
Microsoft Research

ABSTRACT

Software patching has not been effective as a first-line defense against large-scale worm attacks, even when patches have long been available for their corresponding vulnerabilities. Generally, people have been reluctant to patch their systems immediately, because patches are perceived to be unreliable and disruptive to apply. To address this problem, we propose a first-line worm defense in the network stack, using *shields* — vulnerability-specific, exploit-generic network filters installed in end systems once a vulnerability is discovered, but before a patch is applied. These filters examine the incoming or outgoing traffic of vulnerable applications, and correct traffic that exploits vulnerabilities. Shields are less disruptive to install and uninstall, easier to test for bad side effects, and hence more reliable than traditional software patches. Further, shields are resilient to polymorphic or metamorphic variations of exploits [43].

In this paper, we show that this concept is feasible by describing a prototype Shield framework implementation that filters traffic above the transport layer. We have designed a safe and restrictive language to describe vulnerabilities as partial state machines of the vulnerable application. The expressiveness of the language has been verified by encoding the signatures of several known vulnerabilites. Our evaluation provides evidence of Shield's low false positive rate and small impact on application throughput. An examination of a sample set of known vulnerabilities suggests that Shield could be used to prevent exploitation of a substantial fraction of the most dangerous ones.

Categories and Subject Descriptors

D.4.6 [**Security and Protection**]: [Invasive software, Information flow controls]

General Terms

Security, Design, Languages, Performance vspace-0.3cm

Keywords

Worm Defense, Patching, Vulnerability Signature, Network Filter, Generic Protocol Analyzer

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
SIGCOMM'04, Aug. 30–Sept. 3, 2004, Portland, Oregon, USA.
Copyright 2004 ACM 1-58113-862-8/04/0008 ...$5.00.

1. INTRODUCTION

One of the most urgent security problems facing administrators of networked computer systems today is the threat of remote attacks on their systems over the Internet, based on vulnerabilities in their currently running software. Particularly damaging have been self-propagating attacks, or "worms", which exploit one or more vulnerabilities to take control of a host, then use that host to find and attack other hosts with the same vulnerability.

The obvious defense against such attacks is to prevent the attack by repairing the vulnerability before it can be exploited. Typically, software vendors develop and distribute reparative "patches" to their software as soon as possible after learning of a vulnerability. Customers can then install the patch and prevent attacks that exploit the vulnerability.

Experience has shown, however, that administrators often do not install patches until long after they are made available — if at all [33]. As a result, attacks — including worms, such as the widely publicized CodeRed [7], Slammer [40], MSBlast [24], and Sasser [35] worms — that exploit known vulnerabilities, for which patches had been available for quite some time, have nevertheless been quite "successful", causing widespread damage by attacking the large cohort of still-vulnerable hosts. In fact, more than 90% of the attacks today are exploiting known vulnerabilities [1].

There are several reasons why administrators may fail to install software patches:

- *Disruption:* Installing a patch typically involves rebooting, at the very least, a particular host service, and possibly an entire host system. An administrator for whom system and service uptime are crucial may therefore be unable to tolerate the required service or system disruption.

- *Unreliability:* Software patches are typically released as quickly as possible after a vulnerability is discovered, and there is therefore insufficient time to do more than cursory testing of the patch. For popular software programs, patch (regression) testing is inherently difficult and involves an exponential number of test cases due to the software's numerous versions, and their dependencies on various versions of libraries — which depend in turn on other libraries, and so on.

  In addition, a patch typically bundles a number of software updates together with the security fix, further complicating testing. (Such bundling is necessary to prevent the number of branches in the state of the software source code from exploding exponentially.) Hence patches can have serious undetected side

SAMSUNG 1024

effects in particular configurations, causing severe disruption and even damage to the host systems to which they are applied. Rather than risk such damage, administrators may prefer to do their own thorough and time-consuming testing, or simply wait — accepting the risks of vulnerability in the meantime — until the patch has been vindicated by widespread uneventful installation [2].

- *Irreversibility:* Most patches are not designed to be easily reversible due to the ordering of changes that have been made to the system. Once a patch is applied, there is often no easy way of uninstalling it, short of restoring a backup version of the entire patched application (or even the entire system). This factor exacerbates the risk associated with applying a patch.

- *Unawareness:* An administrator may simply miss a patch announcement for some reason, and therefore be unaware of it, or have received the announcement but neglected to act on it.

Because of the above drawbacks to installing patches, methods are being explored for mitigating vulnerabilities without installing patches, or at least until a patch is determined to be safe to install. The goal is to address the window of vulnerability between vulnerability disclosure and software patching. A firewall, for example, can be configured to prevent traffic originating outside a local network from reaching a vulnerable application, by blocking the appropriate port. By doing so, it can protect an application from attacks from "outside". Of course, blocking all traffic to a port is a crude measure, preventing the application from functioning at all across the firewall. Exploit-signature-based firewalls that string-match network traffic to known attack patterns (using regular expressions, for example) are more flexible [41], but metamorphic or polymorphic variations [43] of the known exploits can easily undermine such filters. Furthermore, such exploit signatures can only be obtained after the onset of the attacks themselves, and for fast-spreading worms, it is challenging to detect, extract, and distribute attack signatures in a timely matter [39, 42]. Ideally, prevention mechanisms would be deployed well before attacks occur, and even before the public disclosure of a vulnerability, and only traffic that exploits the vulnerability would be blocked, while all other traffic would be allowed to pass through to the application.

In this paper, we explore the possibility of applying an intermediate "patch" in the network stack to perform this filtering function, to delay (or perhaps in some cases even eliminate) the need for installing the software patch that removes the vulnerability. *Shield* is a system of vulnerability-specific, *exploit-generic* network filters (*shields*) installed at the end host, that examines the incoming or outgoing traffic of vulnerable applications and corrects the traffic according to the vulnerability signature. Shield operates at the application protocol layer, *above* the transport layer. For example, a shield (conceptually, there is one shield per vulnerability) designed to protect against a buffer overrun vulnerability would detect and drop or truncate traffic that resulted in an excessively long value being placed in the vulnerable buffer. Shield differs from previous anti-worm strategies (see Section 9) in attempting to remove a specific vulnerability directly before the vulnerability disclosure time, rather than mitigating or countering the effects of its exploitation at attack time.

Unlike software patches, shields are deployed in the network stack of the end host. They are thus more separated from the vulnerable application — and its wide variety of potential environment configurations — and therefore less likely to have unforeseen side effects. In particular, their compatibility with normal operation is in principle relatively easy to test: Since they operate only on network traffic, and are intended only to drop attack traffic, they can be tested simply by exposing them to a suitably rich collection of network traffic — such as a long trace of past network activity, or a synthetic test suite of representative traffic — to verify that they would allow it all through unaffected. Furthermore, shields are non-disruptive and easily reversible. Automatic installation of shields is therefore a technically viable option.

The most efficient kind of end-host Shield would be positioned at the highest protocol layer, namely the application layer — assuming that programmatic "hooks" into that layer are available for traffic interception and manipulation. This way, any redundant message parsing would be avoided. For example, URLScan [8] is essentially a Shield specific to Microsoft's IIS web server, which uses the IIS ISAPI extension package's hooks for HTTP request interception and manipulation. However, most applications do not offer such extensibility into their protocols. To this end, we have designed a Shield framework that lies between the application layer and the transport layer and offers shielding for *any* application level protocols. Being above the transport layer, Shield does not need to deal with IPSec-encrypted traffic. Encrypted traffic above the transport layer, such as SSL-[13] or application-specific encrypted traffic, are difficult for such a framework to handle[1]. Nevertheless, it is sensible to build the Shield framework to encompass commonly used protocols such as SSL [13] and RPC [34], and the techniques described in this paper can be readily applied to them.

In our Shield framework, we model vulnerability signatures as a combination of partial protocol state machines and vulnerability-parsing instructions for specific payloads (Section 2). For generality, we abstract out the generic elements of application-level protocols into a *generic* protocol analyzer in our Shield architecture (Section 3). For flexibility and simplicity, we express vulnerability signatures and their countermeasures in a safe, restrictive, and yet expressive policy language, interpreted by the Shield framework at runtime (Section 5). We also minimize Shield's maintenance of protocol and message parsing state for scalability (Section 4.1). In particular, Shield reconstructs vulnerability-related application semantics by maintaining the relevant protocol context of a communication session, and performs application-message-based inspection rather than packet-level inspection, as used by some Network Intrusion Detection or Prevention Systems [41]. Packet-level inspection can be easily evaded or disturbed by attackers [32], resulting in both false positives and false negatives. We have implemented a preliminary Shield prototype and experimented with a number of known vulnerabilities, including the ones behind the (in)famous MSBlast, Slammer, and CodeRed worms (Section 7). Our evaluation provides evidence of zero false positives and manageable impact on application throughput (Section 8). An examination of a sample set of known vul-

[1]Nonetheless, assuming Shield runs with root or administrator privileges on end hosts, it would be possible for Shield to obtain encryption keys and decrypt traffic itself. The resulting overhead should not be prohibitive as long as efficient symmetric ciphers are used for the encryption.

nerabilities suggests that Shield could be used to prevent exploitation of a substantial fraction of the most dangerous ones (Section 8.1).

2. VULNERABILITY MODELING AND SHIELD USAGE

An essential part of the Shield design is the method of modeling and expressing vulnerability signatures. A *Shield vulnerability signature* specifies all possible sequences of application messages and specific payload characteristics that lead to *any* remote exploit of the vulnerability. (Note that not all vulnerabilities are suitable for shielding. This issue is discussed further in Section 8.1.) For example, the signature for the vulnerability behind the Slammer worm [40] is the arrival of a UDP packet at port 1434 with a size that exceeds the legal limit of the vulnerable buffer used in the Microsoft SQL server 2000 implementation. More sophisticated vulnerabilities require tracing a sequence of messages leading up to the actual message that can potentially exploit the vulnerability.

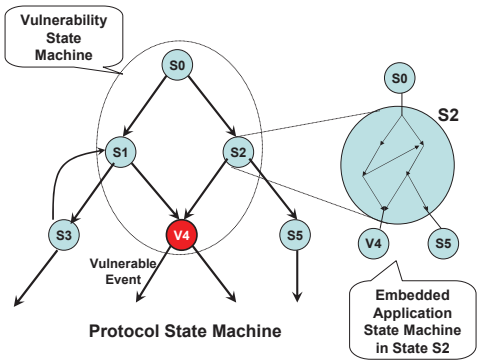

Figure 1: Vulnerability Modeling

To express vulnerability signatures precisely, we have developed a taxonomy for modeling vulnerabilities, as illustrated in Figure 1. Each application can be considered as a finite state machine, which we call the *application state machine*. Overlaying the application state machine is the *Protocol State Machine*, in which the transitions are application message arrivals. The protocol state machine is much smaller and simpler than the application state machine, and the application state machine (e.g., enlarged state $S2$ in the figure) can be viewed as a refinement of the protocol state machine. As a network filter, Shield is primarily concerned with the protocol state machine. We define the *pre-vulnerability state* as the state in the protocol state machine at which receiving an exploitation network event could cause damage ($v4$ in the figure). We call the partial protocol state machine that leads to the pre-vulnerability state the *vulnerability state machine*, and the message that can potentially contain an exploit, the *vulnerable event*.

A Shield vulnerability signature essentially specifies the vulnerability state machine and describes how to recognize any exploits in the vulnerable event. A Shield *policy* for a vulnerability includes both the vulnerability signature and the actions to take on recognizing an exploit of the vulnerability. In Section 5, we detail our design of a language for Shield policy specification.

When a new vulnerability is discovered, a shield designer — typically the vulnerable application's vendor — creates a Shield policy for the vulnerability (conceptually, there is one Shield policy per vulnerability) and distributes it to users running the application as in anti-virus signature distribution and installation. Using this policy, Shield intercepts its application's traffic and walks through the vulnerability state machine; when reaching the pre-vulnerability state, the Shield examines the vulnerable event for possible exploits and takes the specified actions to protect against the exploits if they are present.

Shields can protect a host from potentially malicious incoming traffic in a similar fashion to a firewall, but with much more application-specific knowledge. In addition, shields can filter out *outgoing* traffic that triggers malicious responses back to the host itself, or protects other hosts on the same local network from the host's own vulnerability-exploiting outgoing traffic. The latter use assumes that the shield is installed at a higher privilege level than the malicious or compromised sender of the vulnerability-exploiting traffic. Otherwise, the sender could simply disable the shield before attacking the other hosts.

Shield policies are active as soon as they are received by end host Shield and do not require restarting the vulnerable service or rebooting the system. Once a software patch is applied to the vulnerable application, the corresponding policy can be removed from Shield.

3. SHIELD ARCHITECTURE

3.1 Goals and Overview

The objective of Shield is to emulate the part of the application level protocol state machine that is relevant to its vulnerabilities with intercepted messages and counter any exploits at runtime.

We identify three main goals for the Shield design:

1. *Minimize and limit the amount of state maintained by Shield*: Shield must be designed to resist any resource consumption ("Denial-Of-Service", or "DoS") attacks. Therefore, it must carefully manage its state space. For an end-host-based Shield, the bar is not high: Shield only needs to be as DoS-resilient as the service it is shielding.

2. *Enough flexibility to support any application level protocol*: Flexibility must be designed into Shield so that vulnerabilities related to any application level protocol can be protected by Shield. Moreover, the Shield system design itself should be independent of specific application level protocols, because the Shield system design and implementation would simply not scale if it were necessary to add individual application level protocols to the core system one at a time.

3. *Design Fidelity*: We must design Shield in such a way that Shield does not become an easier alternative attack target. A robust Shield design must ensure that Shield's state machine emulation is consistent with the actual state machine running in the vulnerable application under all conditions. In other words, it is crucial for us to defend against carefully crafted malicious messages that may lead to Shield's misinterpretation of the application's semantics.

Shield achieves goal 2 by applying the well-known principle of "separating policy from mechanism". Shield's mechanism is generic, implementing operations common among all application level protocols. Shield policies specify the varying aspects of individual application level protocol design as

well as the corresponding vulnerabilities. This separation enables the flexibility of Shield to adapt to various application level protocols.

We identify the following mechanisms as the necessary generic elements of an application level protocol implementation: (Less obvious generic elements will be explained throughout the next section.)

- Application level protocols between two parties (say, a client and a server) are implemented using finite state automata.
- To carry out state machine transitions, each party must perform event identification and session dispatching. *Session* is the unit of communication that is modeled by the protocol state machine. A session contains one or more application messages and one or more protocol states. As a result, application-level messages must indicate a message type and session ID when the number of states exceeds one.
- Implementations of datagram-based protocols must handle out-of-order application datagrams for their sessions. (See Section 4.2.)
- Implementations of application-level protocols that allow message fragments must handle fragmentation. (See Section 4.3.)

The policy specifies the following:

- *Application identification*: how to identify which packets are destined for which application. The service port number typically serves this purpose.
- *Event identification*: how to extract the message type from a received message.
- *Session identification* (if applicable): how to determine which session a message belongs to.
- *State machine specification*: the states, events, and transitions defining the protocol automaton. In our setting, the specification is for the vulnerability state machine, a subgraph of a complete protocol state machine (see Section 2).

We first present the Shield mechanisms, including the policy-enabling mechansims, in Section 3.2. Then, we present our policy language in Section 5.

3.2 Components and Data Structures

In this section, we describe the essential components and data structures of an end-host Shield system. Figure 2 depicts the Shield architecture.

3.2.1 Data Structures

There are two main data structures: the *application vulnerability state machine specifications* ("Spec") and the runtime *session states*.

The *Policy Loader* transforms Shield policies into Specs. Multiple vulnerability state machines for the same application are compiled into one application vulnerability state machine specification. Therefore, there is effectively one state machine specification or Spec per application. Merging vulnerability state machines is easy if an entire protocol had been specified in our Shield language (Section 5) ahead of time: Newly discovered vulnerabilities can be expressed by annotating the already-specified protocol state machine.

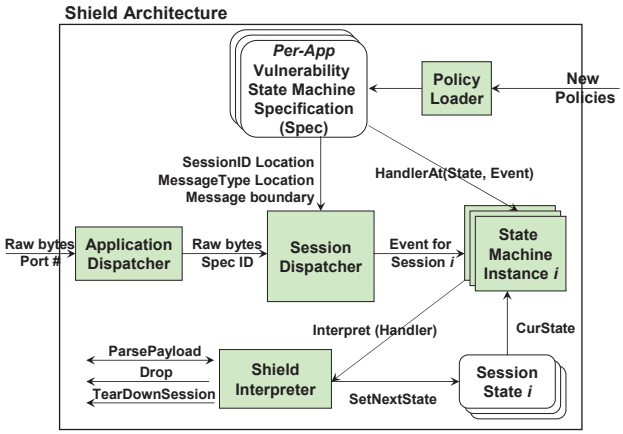

Figure 2: Shield Architecture

The purpose of a Spec is to instruct Shield how to emulate the application vulnerability state machines at runtime. As mentioned in Section 3.1, the Spec contains the state machine specification, port number(s) for application identification, and event and session identification information.

For event and session identification, a Spec indicates the location (i.e., offset and size) vector for the event type and session ID information in the packet, as well as the event type values that are of concern to Shield. For application protocols with only one state, the session ID is unnecessary and left unspecified. Sometimes, an application-level protocol may involve negotiating for a dynamically selected port number as a session ID for further communications (e.g., FTP [30] and RTP [36]). In this case, the port negotation part of the protocol state machine must be specified, and the new port number would be registered with Shield at runtime for application identification from that point on. The session ID would be specified as "PORT", indicating that all communication on this port is considered as a single session. Upon termination of the session, the dynamic port is de-registered with Shield.

To generalize the event dispatching abstraction to text-based application level protocols such as HTTP [12] and SMTP [17], we allow the units of "offset" and "size" to be defined as *words* (made of characters), in addition to bytes. For example, in HTTP and SMTP, the message type is indicated at offset 0, with a size of 1 word. In HTTP, this field contains the request-line method, such as "GET" or "POST"; when it is an HTTP version, it represents a status message type [12]. In SMTP, this field contains the SMTP command, such as "MAIL", "RCPT", or "DATA". (Please see Figure 4 for example.) Using words as a unit requires us to include an additional policy element: maximum word size — otherwise attackers could attack Shield using extremely long words. Of course, we could generalize the unit even further by making unit delimiters and maximum unit size part of the Spec, beyond bytes and space-delimited words. However, we have not yet found this to be necessary for the handful of protocols we have examined.

Some application-level protocols (such as HTTP) allow multiple application-level messages to be received in a single buffer. Therefore, in addition to the session ID and message type, the Spec also specifies the *application level message boundary marker*, if any. For example, for HTTP, the message boundary marker is CRLF CRLF.

One key challenge is that application level messages may

not be received in their entirety (due to congestion control or application-specific socket usage) or in order (due to using UDP, for example). Even the essential event-identifying parts of a message, such as event type and session ID, may not arrive together. We address this problem using DoS-resilient copying or buffering, which is detailed in Section 4.

Note that the *session* is an important abstraction for packet dispatching and as a unit of shielding, apart from socket descriptors or host pairs. This is because one socket descriptor may be used for multiple sequential sessions; and multiple sockets may be used to carry out communications over one session (e.g., FTP [30]). Similarly, one pair of hosts may be carrying out multiple sessions. In these cases, the use of sessions eliminates any ambiguities on which packets belong to which session.

The other data structure in Shield is *session state*. At runtime, Shield maintains session state for each potentially vulnerable communication session. The session state includes the current state of the session and other context information needed for shielding.

3.2.2 Shield Modules

We now describe each Shield module in turn:

Policy Loader: Whenever a new Shield policy arrives or an old policy is modified, the Policy Loader integrates the new policy with an existing Spec if one exists, or creates a new one otherwise. The Shield policy is expressed in the Shield policy language. Policy loading involves syntax parsing, and the resulting syntax tree is also stored in the Spec for the purpose of run-time interpretation of shielding actions For details on the policy language design and interpretation, please see Section 5.

Application Dispatcher: When raw bytes arrive at Shield from a port, the *Application Dispatcher* is invoked to determine which Spec to reference for the arrived data, based on the port number. The Application Dispatcher forwards the raw bytes and the identified Spec to the Session Dispatcher for event and session identification.

Session Dispatcher: On obtaining the locations of the session ID, message type, and message boundary marker from the corresponding Spec, the Session Dispatcher extracts multiple messages (if applicable), recognizes the event type and session ID, and then dispatches the event to the corresponding state machine instance.

State Machine Instance (SMI): There is one state machine instance per session. Given a newly-arrived event and the current state maintained by the corresponding session state, the SMI consults the Spec regarding which event handler to invoke. (Event handlers are included in Shield policies; please see Section 5.) Then the SMI calls the Shield Interpreter to interpet the event handler.

Shield Interpreter: The Shield Interpreter interprets the event handler, which specifies how to parse the application-level protocol payload and examine it for exploits. It also carries out actions like packet-dropping, session tear-down, registering a newly-negotiated dynamic port with Shield, or setting the next state for the current SMI.

4. DETAILED DESIGN ISSUES

4.1 Scattered Arrivals

Although Shield intercepts traffic above the transport layer and does not need to cope with network-layer fragments, each data arrival perceived by Shield does *not* necessarily represent a complete application level message that is independently interpretable by the application. The scattered arrivals of a single application level message could be due to TCP congestion control or some specific message-handling implementations of an application. For instance, a UDP Server may make multiple calls to $recvfrom()$ to receive a single application level message. In this case, Shield would recognize multiple data arrivals for such messages. This complicates session dispatching when session ID or message type are not received "in one shot". It also complicates payload parsing in the event handlers when not enough data has arrived for an event-handler to finish parsing and checking. Here, we must *copy* (i.e., buffer and pass on) *part of* the incompletely-arrived data in the Shield system, and wait for the rest of the data to arrive, before we can interpret it.

In addition, we need to *index* copy-buffers so that later arrivals of the same message for a given session can be stitched together properly. Although socket descriptors are not appropriate to identify sessions (Section 3.2), they are safe for indexing copy-buffers: While multiple sockets could be used for one session (e.g., FTP [30]), a single application message is typically[2] not scattered over multiple sockets — otherwise the application would not be able to interpret the parts of the message due to the lack of information such as session ID or message type; similarly, although one socket descriptor could be used for multiple parallel sessions, an application message has to be received on a socket continuously to its completion without interruptions from any other application messages. Therefore, we can apply *per-socket* copying for incomplete message arrivals.

We differentiate between *pre-session copying* and *in-session copying*. *Pre-session copying* happens when the session ID information has not completely arrived, while *in-session copying* refers to the copying of the data whose session is known. A copy-buffer is associated with a socket initially before a session ID fully arrives. Once the session ID is received, the copy-buffer is associated with both a socket and its respective session. Once a complete application message has been received, the copy-buffer is de-allocated.

We do *not* need to save the entire partially arrived message, but only the partially arrived *field*. For example, when a Session ID field has not arrived completely — say only 2 out of 4 bytes — Shield only needs to remember that it is parsing the Session ID field, and saves the received two bytes only. This reduces the overhead of copying in Shield.

We now introduce another runtime data structure that needs to be maintained by Shield: *parsing state*. This state is per application-level message, and it records which field of an application message is being parsed, and how many bytes have already been received for that field. The field has to be a "terminal" field rather than a structure of fields; for a field nested in other structures or an array, the field is represented as something like "someStructure.fields[i]". This restriction minimizes the amount of copying needed, since terminal fields are typically small.

[2] pTCP [15] proposes the use of TCP-v as an abstraction of a connection which can use multiple sockets instead of one as in TCP. If pTCP were deployed, Shield would use TCP-v to index the copy-buffer instead.

For a vulnerable application, we must maintain the state of the current field being parsed for *each* of the application messages, *even* when Shield had already determined that the session to which the message belongs would not lead to any exploit. (However, we do not maintain session state and the copy-buffer for such sessions.) This prevents ambiguity: If we do not keep the parsing state for the message, other parts of the message would be treated as new application messages. Attackers could then easily craft parts of a single application-level message [32], send them separately, and cause inconsistencies between the emulated state machine in Shield and the actual state machine in the application.

For Shield to be able to parse application messages, parsing instructions (or payload formats) for all of an application's message types must be specified to Shield in policy descriptions. Therefore, payload formats must also be part of the Spec. Fortunately, Shield does not need to parse all messages in detail, but only the parts necessary for detecting the presence of an exploit. Therefore, we can aggressively bundle many fields of an application message into one field with a total byte count or word count, which will be the number of bytes or words to skip during parsing. When Shield determines the innocence of a session, the goal of parsing its subsequent application-level messages is only to find the end of those messages. Hence, parsing for those messages can be even further streamlined.

While specifying all application messages seems daunting, if an application level protocol were specified in a standard and formalized format (such as our policy language-like format — see Section 5), we could automatically extract payload format specifications and vulnerability state machines from that format to our policy language. On the other hand, if a Shield designer knows for a fact that scattered arrivals of a message do never happen (e.g., single rather than multiple $recvfrom()$ calls when receiving a single application message in a UDP server implementation), then only events involved in the vulnerability state machine need to be specified.

4.2 Out-of-Order Arrivals

When an application-level protocol runs on top of UDP, its datagrams can arrive out of order. Applications that care about the ordering of these datagrams will have a sequence number field in their application-level protocol headers. For Shield to properly carry out its exploit detection functions, Shield copies the out-of-order datagrams, and passes them on to the applications. This way, Shield can examine the packets in their intended sequence. Shield sets the upper limit of the number of copied datagrams to be the maximum number of out-of-order datagrams that the application-level protocol can handle. Hence, this maximum also needs to be expressed in the policy descriptions, as does the sequence number location.

4.3 Application Level Fragmentation

Shield runs on top of the transport layer. Hence, Shield does not need to deal with network-layer fragmentation and re-assembly. Nonetheless, some application-level protocols use application data units and perform application level fragmentation and re-assembly. For protocols over TCP, bytes are received in order. For protocols over UDP, Shield copies their out-of-order datagrams to retain the correct packet sequence (see the above section). Therefore receiving and processing application level fragments is no different from processing partially arrived data, as explained in Section 4.1. However, the Spec needs to contain the location of the application-level fragment ID in the message, so that a fragment is not treated as an entire message event.

5. SHIELD POLICY LANGUAGE

In this section, we present the Shield policy language which is used to describe the vulnerabilities and their countermeasures for an application.

Figures 3, 4, and 5 show some examples of our policy language usage. They are policy scripts for the vulnerabilities behind Slammer [40], CodeRed [7], and MSBlast [24], respectively. Please note that these scripts are written based on the knowledge of the vulnerabilities rather than their respective attack instances.

```
# Vulnerability behind Slammer
SHIELD (Vulnerability_Behind_Slammer, UDP, (1434))

# 0 offset, size of 1 byte
MSG_TYPE_LOCATION = (0, 1);

INITIAL_STATE S_WaitForSSRPRequest;
FINAL_STATE S_Final;

# MsgType = 0x4
EVENT E_SSRP_Request = (0x4, INCOMING);

STATE_MACHINE = {
(S_WaitForSSRPRequest, E_SSRP_Request, H_SSRP_Request),
};


HANDLER H_SSRP_Request (DONT_CARE) {
  COUNTER legalLimit = 128;
   # MSG_LEN returns legalLimit + 1 when legalLimit is exceeded
  COUNTER c = MSG_LEN (legalLimit);
  IF (c > legalLimit)
    DROP;
    RETURN (S_FINAL);
  FI
  RETURN (S_Final);
};
```

Figure 3: Policy description of the vulnerability behind Slammer

```
# Shield for vulnerability behind CodeRed
SHIELD(Vulnerability_Behind_CodeRed, TCP, (80))

INITIAL_STATE       S_WaitForGetRequest;
FINAL_STATE         S_Final;

#
MSG_TYPE_LOCATION= (0, 1) WORD;

MSG_BOUNDARY = "\r\n\r\n";

EVENT E_GET_REQUEST = ("GET", INCOMING);

STATE_MACHINE = {
(S_WaitForGetRequest, E_GET_Request, H_Get_Request),
};

PAYLOAD_STRUCT {
  WORDS(1)    method,
  WORDS(1)    URI,
  BYTES(REST) dummy2,
} P_Get_Request;

HANDLER H_Get_Request (P_Get_Request) {
   COUNTER legalLimit = 239;
   COUNTER c = 0;

   # \?(.*)$ is the regular expression to retrieve the
   # query string in the URI
   # MATCH_STR_LEN returns legalLimit + 1 when legalLimit is exceeded
   c = MATCH_STR_LEN (>>P_Get_Request.URI, "\?(.*)$", legalLimit);
   IF (c > legalLimit)
      # Exploit!
      TEARDOWN_SESSION;
      RETURN (S_FINAL);
   FI
   RETURN (S_FINAL);
};
```

Figure 4: Policy description of the vulnerability behind CodeRed

There are two parts to the policy specification in the Shield language. The first part is static and includes states,

```
# SHIELD (Name, Transport_Protocol, (port-list))
SHIELD (Vulnerability_Behind_MSBlast, TCP, (135, 139, 445))

SESSION_ID_LOCATION = (12, 4); # offset 12, 4 bytes
MSG_TYPE_LOCATION = (2, 1);    # offset 2, 1 byte

INITIAL_STATE S_WaitForRPCBind;
FINAL_STATE   S_Final;
STATE       S_WaitForRPCBindAck;
STATE       S_WaitForRPCAlterContextResponse;
STATE       S_WaitForRPCRequest;
STATE       S_WaitForSessionTearDown;

# EVENT eventName = (<eventTypeValue>, <direction>)
EVENT E_RPCBind         = (0x0B, INCOMING);
EVENT E_RPCBindAck = (0x0C, OUTGOING);
EVENT E_RPCBindNak = (0x0D, OUTGOING);
EVENT E_RPCAlterContext = (0x0E, INCOMING);
EVENT E_RPCAlterContextResponse = (0x0F, OUTGOING);
EVENT E_RPCRequest = (0x0,  INCOMING);
EVENT E_RPCShutdown = (0x11, OUTGOING);
EVENT E_RPCCancel = (0x12, INCOMING);
EVENT E_RPCOrphaned = (0x13, INCOMING);

STATE_MACHINE = {
# (State, Event, Handler),
(S_WaitForRPCBind, E_RPCBind, H_RPCBind),
(S_WaitForRPCBindAck, E_RPCBindAck, H_RPCBindAck),
(S_WaitForRPCRequest, E_RPCRequest, H_RPCRequest),
...
};

# payload parsing instruction for P_Context
PAYLOAD_STRUCT {
   SKIP BYTES(6)                        dummy1,
   BYTES(1)                 numTransferContexts,
   SKIP BYTES(1)                        dummy2,
   BYTES(16)                 UUID_RemoteActivation,
   SKIP BYTES(4)                        version,
   SKIP BYTES(numTransferContexts * 20) transferContexts,
} P_Context;
```

```
PAYLOAD_STRUCT {
   SKIP BYTES(24)              dummy1,
   BYTES(1)                    numContexts,
   SKIP BYTES(3)        dummy2,
   P_Context[numContexts]      contexts,
   ...
} P_RPCBind;


HANDLER H_S_RPCBind (P_RPCBind)
{
   # if invoking the RemoteActivation RPC call
   IF (>>P_RPCBind.contexts[0]
== 0xB84A9F4D1C7DCF11861E0020AF6E7C57)
      RETURN (S_WaitForRPCBindAck);
   FI
   RETURN (S_Final);
};

HANDLER H_RPCBindAck (P_RPCBindAck)
{
  RETURN (S_WaitForRPCRequest);
};

HANDLER H_RPCRequest (P_RPCRequest)
{
   IF (>>P_RPCRequest.bufferSize > 1023)
      TEARDOWN_SESSION;
      PRINT ("MSBlast!");
      RETURN (S_Final);
   FI
   RETURN (S_WaitForSessionTearDown);
};

# other PAYLOAD_STRUCTs/Handlers
# are  not included here ...
```

Figure 5: Excerpt from the policy description of the vulnerability behind MSBlast

events, state machine transitions, and generic application level protocol information such as ports used, the locations of the event type, session ID, sequence number or fragment ID in a packet, and the message boundary marker. This part of the policy specification is loaded into the Spec data structure directly by the Policy Loader (Figure 2), and is independent of runtime conditions.

The second part of the policy specification is for runtime interpretation during exploit checking. This includes the handler specification and payload parsing instructions (i.e., PAYLOAD_STRUCT definitions in the figures). The role of the handler is to examine the packet payload and pinpoint any exploit in the current packet payload, or to record the session context that is needed for a later determinination of exploit occurrence. To examine a packet, a handler needs to follow the policy's payload parsing instructions.

When a policy is loaded, the Policy Loader parses the syntax of the handlers and the payload format, and stores the syntax tree in the Spec for run-time interpretation.

5.1 Payload Specification

The PAYLOAD_STRUCT definitions specify how to parse an application-level message. Shield needs not parse out all the fields of a payload, as in the actual applications, but only the fields relevant to the vulnerability. We allow the policy writers to simplify payload parsing by clustering insignificant fields together as a single dummy field of the required number of bytes (e.g., field dummy1 of P_RPC_Bind in Figure 5). Such fields are marked as skippable during parsing using the keyword SKIP, so that no copy-buffer is maintained for such fields (Section 4.1). From examining a number of application level protocols, we have found that payload parsing specification only needs to support a limited set of types for fields, including bytes of any size (i.e., BYTES(num) where "num" could be a variable size or an expression), words of any size (i.e., WORDS(num)) for text-based protocols, (multi-dimensional) arrays of PAYLOAD_STRUCT's, and booleans.

5.2 Handler Specification

The Shield language for handler specification is very simple, and highly specialized for our purpose. Variables have only two scopes: they are either local to a handler or "global" within a session across its handlers. There are only four data types: BOOL for booleans, COUNTER for whole numbers, byte arrays such as "BYTES(numBytes)", and word arrays such as "WORD(numWords)". We also have built-in variables for handlers to use, such as SESSION_ID.

Built-in functions include DROP , TEARDOWN_SESSION, REGISTER_PORT, length-based functions such as MSG_LEN (Figure 3) and MATCH_STR_LEN (Figure 4), and regular expression functions that may be needed for text-based protocols. DROP drops a packet while TEARDOWN_SESSION closes all sockets associated with a session. The regular expression functions are data stream-based rather than string buffer-based. That is, they must be able to cope with scattered message arrivals. Similarly, length functions are also stream-based with a required parameter of "stop count" (e.g., "legalLimit" in our example scripts) to facilitate the handling of buffer overrun-type vulnerabilities. When the length reaches "stop count", counting stops and returns "stop count"+1. This way, Shield will not count and maintain state beyond what is necessary in the case of buffer-overrun exploits.

The syntax "≫payload" instructs Shield to parse and to refer to the bytes that represent the "payload" of the packet, according to the parsing instruction defined for "payload" (i.e., there should be a definition: "PAYLOAD_STRUCT {...} payload;" earlier in the policy decription). The parsed fields of a PAYLOAD_STRUCT are treated as local variables for that handler. Within the handler, we allow assignments, if-statements, iterators, and return-statements that exit the handler and indicate the next state the session should be in. The iterators, rather than being traditional general-purpose for-loops, are used for parsing iterative payload structures, such as arrays of items. For example, given "FOREACH

(item IN ≫ Payload.itemArray) { ... }", the interpreter parses items of the "itemArray" field of the "Payload" iteratively, according to the "Payload" definition, and along the way, performs some operations on the bytes representing each "item". Note that the interpreter does not keep state across items.

During handler interpretation, the current payload being parsed may not be completely received. In this case, we save the execution state of the handler as part of the session state so that when new data arrives, the handler's execution can be resumed. This is very much like call continuation. In our case, the continuation state includes a queue of current handler statements being executed because of potentially nested statements and the parsing state (Section 4.1) for the payload — that is, the current field of the payload being parsed, and the bytes read for that field.

The restrictive nature of our language makes it a safer language than general-purpose languages. While our language is restrictive, we find it sufficient for all of the vulnerabilities that we have worked with and the application level protocols that we have examined. However, it is still evolving as we gain experience from shielding more vulnerabilities and protocols; and in the process, we are also investigating how we may incorporate techniques from previous research in protocol, packet, or data structure specification languages [21, 25, 11], that are appropriate for the purphose of shielding.

Our language is simpler than the Bro language [27] that is used for scripting the security policies of the Bro Nework Intrusion Detection System (NIDS). This is because Bro performs network monitoring and intrusion detection for the network layer and above of the network stack, and also monitors cross-application, cross-session interactions based on various attack patterns. In contrast, Shield is only concerned with application-specific traffic passing over transport layer, or of even higher-level protocols such as HTTP and RPC. Furthermore, a key advantage of the vulnerability-driven approach of Shield over attack- or exploit-driven approachs such as NIDS is that Shield does not need to consider attack activities before the vulnerable application is involved, as in, for example, multi-stage attacks, where an attack setup stage precedes the actual vulnerability exploitation. Shield only needs to screen the traffic of particular vulnerable applications.

6. ANALYSIS

6.1 Scalability with Number of Vulnerabilties

In this section, we discuss how Shield scales with the number of vulnerabilities on a machine.

The number of shields on an end host should *not* grow arbitrarily large, because each shield will presumably be removed when its corresponding vulnerability is patched. Also, Shields are application-specific, adding negligible overhead to applications to which they do not apply. Hence, N Shields for N different applications are equivalent to a single shield in terms of their effect on the performance of any single application.

An application may have multiple vulnerabilities over time. The state machines that model these vulnerabilities should preferably be merged into a single one. Otherwise, each state machine must be traversed for each packet, resulting in linear overhead. When these vulnerabilities appear on disjoint paths of the merged state machine, per-packet shield processing overhead for them is almost equivalent to the overhead for just one vulnerability. For vulnerabilities that share paths in the state machine, however, shield overhead may be cumulative. On the other hand, our data on vulnerabilities presented in Section 8.1 suggests that this cumulative effect is not significant: For worm-exploitable vulnerabilities, no more than three vulnerabilities ever appeared over a single application level protocol throughout the whole year.

In any case, for vulnerable applications, the application throughput with shield is, at worst, about halved, since the network traffic is processed at most twice — once in Shield and once in the application. Moreover, our experiments with our Shield prototype indicate that Shield's impact on application throughput is small (Section 8.2).

6.2 False Positives

By design, shields are able to recognize and filter *only* traffic that exploits a specific vulnerability, and hence should have very low false positives. However, false positives may arise from incorrect policy specification due to misunderstanding of the protocol state machines or payload formats. Such incorrect policy specification can be debugged with stress test suites or simply by replaying a substantial volume of application traffic traces. Trace replay at the application level is easy since it is not necessary to replay the precise transport protocol behavior.

Another source of false positives may be state-sensitive application behavior upon receiving an exploit event. The application state machine embedded at that state may only trigger the vulnerable code based on the local machine setting or some runtime conditions. While such information can also be incorporated into Shield, it is difficult to generalize such application-specific implementation details to simple and safe policy language constructs. For the vulnerabilities with which we have experimented, we have not observed such false positives.

7. IMPLEMENTATION

We have prototyped an end host-based Shield system on Microsoft Windows XP. In particular, we have implemented Shield as a Microsoft WinSock2 Layered Service Provider (LSP) [16]. WinSock2 API is the latest socket programming interface for network applications on Windows. At runtime, these network applications link in the appropriate socket functions from the WinSock2 dymamically linked library (DLL) when making socket function calls. The LSP mechanism in WinSock2 allows new service providers to be created for intercepting WinSock2 calls to the kernel socket system calls. An LSP is compiled into a dynamically linked library. Upon installation, any applications making WinSock2 calls link in both the WinSock2 DLL and the LSP DLL. We use this mechanism to implement Shield for intercepting vulnerable application traffic above the transport layer as shown in Figure 6.

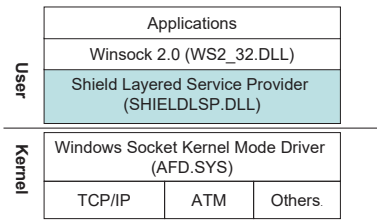

Figure 6: Shield Prototype Using WinSock2 LSP

Our Shield LSP implements the architecture depicted in

Figure 2 with 10,702 lines of C++ code[3]. We employ Flex [26] and Byacc [4] to parse the syntax of the Shield policy language, and the Policy Loader calls the Byacc API to obtain the syntax trees of the policy scripts.

We have used the vulnerabilities behind Slammer [40], MSBlast [24], CodeRed [7], and twelve other vulnerabilities from Microsoft security bulletins to drive our design and implementation. They are all input validation failure vulnerabilities, such as buffer overruns, integer overflows, or malformed URLs. Slammer exploits a proprietary application level protocol, SSRP [40], on top of UDP. MSBlast exploits RPC [34] over either TCP or UDP. CodeRed uses HTTP [12]. Other vulnerabilities exploit Telnet [29], SMB [38], HTTP or RPC. We have also examined some other protocols such as RTP [36] and SMTP [17] during the design of our policy language.

8. EVALUATIONS

8.1 Applicability of Shield

How applicable is Shield to real-world vulnerabilities? Shield was designed to catch exploits in a wide variety of application-level protocols, but there are several potential gaps in its coverage:

- Vulnerabilities that result from bugs that are deeply embedded in the application's logic are difficult for Shield to defend against without replicating that application logic in the network. For example, browser-based vulnerabilities that can be exploited using HTML scripting languages are difficult for Shield to prevent, since those languages are so flexible that incoming scripts would likely have to be parsed and run in simulation to discover if they are in fact exploits.
- Even simple vulnerabilities that are exploitable by malformed, *network protocol-independent* application objects (such as files) are difficult for Shield to catch. For example, a shield against otherwise simple buffer overruns in application file format parsers would have to spot an incoming file arriving over many different protocols. For file-based vulnerabilities, vulnerability-specific anti-virus software (rather than the exploit-signature-based kind typically used today) would be more appropriate. Such software already has to deal with the polymorphism sometimes exhibited by viruses.
- Application-specific encryption poses a problem for Shield, as mentioned in Section 1.

To assess the significance of these obstacles, we analyzed the entire list of security bulletins published by the Microsoft Security Response Center (MSRC) for the year 2003. Table 1 summarizes our findings. Of the 49 bulletins, six described vulnerabilities that were purely local, not involving a network in any way. Of the rest, 24 described "user-involved" vulnerabilities, in the sense of requiring local user action on the vulnerable machine — such as navigating to a malicious Website, or opening an emailed application — to trigger. The remaining 19 described server vulnerabilities, in the sense of being possible to trigger via the network, from outside the machine.

The user-involved vulnerabilities generally appear difficult to design shields for. However, none of them are likely to result in self-propagating worms, because they cannot be exploited without some kind of user action upon the browser. For example, seven involved application file formats. Two were email client vulnerabilities, one was a media player vulnerability, and the rest were found in the browser, and hence invoked via HTML or client-side scripting.

[3]This line count does not include the generated Flex and Byacc files.

# of Vuln.	Nature	Wormable	Shieldable
6	Local	No	No
24	User-involved	No	Usually Hard
12	Server buffer overruns	**Yes**	**Easy**
3	Cross-site scripting	No	Hard
3	Server Denial-of-service	No	Varies

Table 1: Applicability of Shield for vulnerabilities of MSRC over the year 2003.

Of the server vulnerabilities, twelve might conceivably be exploitable by worms, under "ideal" conditions — i.e., the server application being very widely deployed in an unprotected, unpatched and unfirewalled configuration. The remainder included three denial-of-service attacks, three "cross-site scripting" attacks, and a potential information disclosure. These are not vulnerable to exploitation by worms.

Of the potentially worm-exploitable vulnerabilities, five involved application level protocols running over HTTP. The rest involved specific application level protocols — typically directly over TCP or UDP — none of which appear inherently incompatible to the Shield approach. Moreover, all twelve were based on buffer overruns, and hence amenable to shielding.

Thus, while many vulnerabilities may not appear to be suitable for the Shield treatment, in fact the most threatening — those prone to exploitation by worms — appear to be disproportionately Shield-compatible.

We also assessed the reliability of the patches associated with our sample set of security bulletins. Of the patches associated with the 49 bulletins, ten (including three repairing potentially worm-exploitable vulnerabilities) were updated at least once following their initial release. Eight of those (including two involving wormable vulnerabilities) were updated to mitigate reported negative side effects of the patch. The others were augmented with extra patches for legacy versions of the product. These side effects would likely have been avoided had Shield been used in place of the patch, since a key advantage of shields over patches is their easy testability (see Section 1).

Finally, with the exception of HTTP-related vulnerabilities, no single application-level protocol exhibited more than RPC's three vulnerabilities during the entire year. Hence, apart from the HTTP port, no port is likely to be burdened with so many combined shields at a given time that the cumulative performance cost of multiple shields becomes an issue separate from the overhead of shielding the port in the first place.

8.2 Application Throughput

To evaluate the impact of Shield on application throughput and CPU usage, we have devised the following experiment: We have multiple clients establishing simultaneous sessions over TCP with a server. The server is a Dell PWS650 with a 3.06 Ghz CPU and 1 GB of RAM. The clients and the server are connected via a 100 Mbps or 1 Gbps Ethernet switch. All computers run Windows XP SP1.

We use the following client-server protocol:

```
SHIELD (ThroughputTest, TCP, (9898))

STATE_MACHINE = {
(S_WaitForBind, E_Bind, H_Bind),
(S_WaitForBindAck, E_BindAck, H_BindAck),
(S_WaitForRequest, E_Request, H_Request),
(S_WaitForResponse, E_Response, H_Response),
(S_WaitForRequest, E_Shutdown, H_Shutdown),
};

PAYLOAD_STRUCT {
   BYTES(1024) item,
} P_Unit;

PAYLOAD_STRUCT {
   BYTES(4)    field1,
   # upto 1024 bytes of data is buffered for each session
   P_Unit[1024]  MBytes,
}  P_RESPONSE;

HANDLER H_Response (P_RESPONSE)
{
   FOREACH (P_unit IN >>P_RESPONSE.MBytes) {
       # touching each byte of the 1 MB data
       PRINT (>>P_Unit.item);
   }
   RETURN (S_FINAL);
};
```

Figure 7: Excerpt from our throughput experiment policy

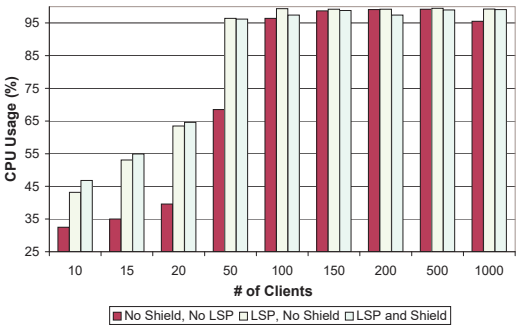

Figure 8: CPU Comparison for 100 Mbps Switch

```
1. Client -> Server: BIND
2. Server -> Client: BIND_ACK
3. Client -> Server: MSG_REQUEST
4. Server -> Client: MSG_RESPONSE of 1 MB data
5. Goto 1
```

We used the policy script shown in Figure 7. 1 MB of data in MSG_RESPONSE is represented as an array of P_Units that consists of 1024 bytes. This increases the experiment's stress on resources, since each session will buffer up to 1KB.

We measure the server's throughput and CPU usage in three scenarios, using either a 100 Mbps Ethernet switch or a Gbps one: 1) without LSP and Shield; 2) with LSP, but without Shield (i.e., the packets just pass through the LSP); 3) with LSP and Shield.

With the 100 Mbps switch, all three scenarios achieve a maximum possible throughput of 92.8 Mbps. However, we observe that LSP incurs 11-28% CPU overhead when the number of clients is low (≤ 50), and Shield logic adds just a few percent on top of the LSP overhead. Figure 8 shows the CPU usage comparison. With a 1 Gbps Ethernet switch, the high-speed switching saturated CPU usage for all cases. Nevertheless, we can observe the differences in throughput for all three scenarios, as shown in Figure 9. As the number of clients increases, the throughput decreases for all scenarios: LSP degrades throughput by 12%, and Shield further degrades the throughput by another 11%. While such overhead is manageable, we observe that much of the overhead is due to the WinSock LSP design [16]. WinSock LSP is

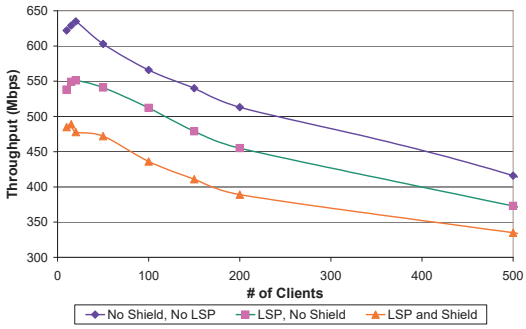

Figure 9: Throughput Comparison for 1 Gbps Switch

designed to allow multiple LSPs (like Shield) to be chained as layers of providers and consumers. A socket instance of an LSP inherits the context information from its provider through socket replication. The base provider is the respective kernel socket. Further, cross-layer socket associations and translations must also be carried out by LSPs. This becomes especially inefficient when "select" calls are used over a large number of sockets for I/O multiplexing, which is the case for our experiment. We suspect that a well-designed kernel implementation of Shield could eliminate much of the overhead incurred by LSP.

8.3 False Positives

As mentioned in Section 6.2, false positives come from either a misunderstanding of the protocol state machine or the differential treatment of an exploit in the application. In this section, we evaluate the false positve nature of our Shield-LSP implementation. We focus our attention on the Shield we designed for Slammer [40], which exploits the SSRP protocol of SQL Server 2000.

We obtained a stress test suite for SSRP from the vendor. SSRP is a very simple protocol with only 12 message types. The test suite contains a total of 36 test cases for exhaustive testing of SSRP requests of various forms. Running this test suite against our Shield, we did not observe any false positives. Of course, this example does not prove that Shield is false positive-free.

9. RELATED WORK

The onsets of CodeRed [7], Slammer [40], and MSBlast [24] in the past few years have provoked great interest in worm defense research. Many [23, 42, 22, 45] have characterized and analyzed the fast- and wide-spreading nature and potential [47, 6] of modern-day worms.

Most worm attacks today exploit software defects, such as buffer overruns on stacks or heaps, for remote code execution. Static checkers that perform data flow analysis [3, 44] over source code have been effective in finding software defects systematically before software releases. However, such tools are not false negative-free. Mitigation techniques, such as Stackguard [9] and non-executable stacks or heaps, raise the bar for software defect exploitation, but attackers can cross the bar with new exploitation techniques [28, 49, 19].

Attack prevention techniques address attacks against known vulnerabilities. These techniques make use of well-defined invariants, namely the known vulnerabilities. Both software patching and Shield fall into this category. Section 1 gives detailed comparisons between the two.

Worm containment techniques are typically used for containing *known* worms. Firewalls can be used for this purpose, by, for example, blocking a port that is under attack. Firewalls have a function similar to Shield, but work in a much cruder way — rarely customized in response to a particular vulnerability, for instance. They are usually not deployed on the end host, and thus vulnerable to evasion opportunities [32] for attackers. Moreover, firewalls are unaware of most application-level protocols (and may not even have access to them — if, for example, traffic is encrypted).

Exploit signature-based Network Intrusion Prevention systems (NIPS), such as Snort-based Hogwash [41], filter out malicious traffic according to attack signatures. The signatures are typically in the form of regular expressions. Traffic blocking is based on packet-level inspection and pattern matching, which can be easily manipulated by attackers to cause false negatives and false positives, since application messages can be scattered over multiple packets [32]. These systems require fast signature extraction algorithms (e.g., EarlyBird [39], Autograph [5]) for them to be worm-containing. The biggest challenge these algorithms face is polymorphic or metamorphic worms.

Rate-limiting [48] is another containment method that throttles the sending rate at an infected end host. The containment method — blocking the detected scanning activity of compromised nodes [46] — has also been explored as a weapon against scanning worms. Rate-limiting and scanning worm containment are mostly useful for fast-spreading worms, rather than the "stealthy" kind.

Network Intrusion Detection Systems (NIDS), exemplified by Bro [27] and Snort [41], monitor network traffic and detect attacks of known *exploits*. NIDS are usually more customized by application than firewalls, but deal with known exploits rather than known vulnerabilities. Unlike Shield, NIDS are not on the traffic forwarding path. Moreover, they focus on detection rather than attack prevention. For more reliable attack detection, "traffic normalizers" [37, 14] or "protocol scrubbers" [20] have been proposed to protect the forwarding path by eliminating potential ambiguities before the traffic is seen by the monitor, thus removing evasion opportunities. In addition to evasion, NIDS systems face the issue of a high false postive rate, which complicates the reaction process.

Another interesting attack detection and signature extraction mechanism is the deployment of "honeypots" that cover "dark" or unused IP address space. Some examples are Backscatter [23], *honeyd* [31], HoneyComb [18], and HoneyStat [10]. Any unsolicited outgoing traffic from the honeypots reveals the ocurrence of some attack.

10. CONCLUSIONS AND FURTHER WORK

We have shown that network-based vulnerability-specific filters are feasible to implement, with low false positive rates, manageable scalability, and broad applicability across protocols. There are, however, still a number of natural directions for future research on Shield:

- Further experience writing shields for specific vulnerabilities will better indicate the range of Shield's applicability and the adequacy of the Shield policy language. It may also be possible to develop automated tools to ease Shield policy generation.

  For example, writing a shield policy currently requires a fairly deep understanding of the protocol over which the vulnerability is exploited. For protocols described in a standard, formalized format, however, it should be possible to build an automated tool that generates most of the protocol-parsing portion of a shield policy. The rest of the task of writing the policy would still be manual, but it is often relatively easy, since the vulnerability-exploiting portion of the incoming traffic — say, an overly long field that causes a buffer overrun — is often easy to identify once the traffic has been parsed.

- Shield needs not necessarily be implemented at the end host. It may be preferable in some cases, from an administration or performance point of view, to deploy Shield in a firewall or router, or even in a special-purpose box. However, these alternate deployment options have yet to be explored.

- One of the advantages of Shield is that shields can in principle be tested in a relatively simple way, verifying that some collection of traffic (test suites or real-world traces) is not interfered with. Automating this process would make the shield release process even easier.

- Ensuring the secure, reliable and expeditious distribution of Shields is crucial. While releasing a patch enables attackers to reverse-engineer the patch to understand its corresponding vulnerability, and thus to exploit it, Shield makes reverse-engineering even easier since vulnerability signatures are spelled out in Shield policies. Therefore, Shield distribution and installation is in an even tighter race with the exploit-designing hacker.

- It is possible that Shield's design might prove useful when applied to the virus problem, since some viruses exploit a vulnerability in the application that is invoked when an infected file is opened. Today, most anti-virus software is signature-based, identifying specific exploits rather than vulnerabilities. Incorporating shield-like technology into anti-virus systems might allow them to protect against generic classes of viruses that use a particular infection method.

11. ACKNOWLEDGEMENT

Jon Pincus has given us insightful and constant advice since the idea formation stage of the Shield project. Jay Lorch gave us many thoughtful critiques on the first draft of this paper. Many Microsoft employees from the product side have graciously helped us with understanding various aspects of many vulnerabilities from Microsoft Security Bulletins and using stress test suites for a number of application level protocols. Andrew Begel and Zhe Yang offered us helpful discussions on our policy language design and interpreter implementation. This work also benefited from our discussions with Nikita Borisov, David Brumley, Hao Chen, John Dunagan, Jason Garms, Jon Howell, Yih-Chun Hu, Jitu Padhye, Vern Paxson, Stefan Savage, Dawn Song, Nick Weaver, and Brian Zill. The final version of this paper is much influenced by the anonymous SIGCOMM reviewers, Nikita Borisov, and our shepherd, Paul Barford. We are grateful for everyone's help.

12. REFERENCES

[1] W. A. Arbaugh, W. L. Fithen, and J. McHugh. Windows of Vulnerability: a Case Study Analysis. *IEEE Computer*, 2000.

[2] Steve Beattie, Seth Arnold, Crispin Cowan, Perry Wagle, and Chris Wright. Timing the application of security patches for optimal uptime. In *LISA XVI*, November 2002.

[3] William Bush, Jonathan D. Pincus, and David J. Sielaff. A Static Analyzer for Finding Dynamic Programming Errors. *Software-Practice and Experience (SP&E)*, 2000.

[4] Byacc. http://dickey.his.com/byacc/byacc.html.

[5] H. Chen and B. Karp. Autograph: Toward automated, distributed worm signature detection. In *Proceedings of the 13th Usenix Security Symposium*, 2004.

[6] Z. Chen, L. Gao, and K. Kwiat. Modeling the Spread of Active Worms. In *Proceedings of IEEE Infocom*, 2003.

[7] Microsoft Security Bulletin MS01-033, November 2003. http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/bulletin/MS01-033.asp.

[8] Microsoft Corp. URLScan Security Tool. http://www.microsoft.com/technet/security/URLScan.asp.

[9] Crispin Cowan, Calton Pu, Dave Maier, Heather Hintony, Jonathan Walpole, Peat Bakke, Steve Beattie, Aaron Grier, Perry Wagle, and Qian Zhang. StackGuard: Automatic Adaptive Detection and Prevention of Buffer-Overflow Attacks. In *Proceedings of 7th USENIX Security Conference*, 1998.

[10] David Dagon, Xinzhou Qin, Guofei Gu, Wenke Lee, Julian Grizzard, John Levine, and Henry Owen. HoneyStat: LocalWorm Detection Using Honeypots. In *RAID*, 2004.

[11] O. Dubuisson. *ASN.1 - Communication Between Heterogeneous Systems*. Morgan Kaufmann Publishers, 2000.

[12] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, and T. Berners-Lee. *Hypertext Transfer Protocol – HTTP/1.1 (RFC 2616)*, June 1999.

[13] Alan O. Freier, Philip Karlton, and Paul C. Kocher. The SSL Protocol Version 3.0. http://wp.netscape.com/eng/ssl3/ssl-toc.html.

[14] Mark Handley, Vern Paxson, and Christian Kreibich. Network Intrusion Detection: Evasion, Traffic Normalization, and End-to-End Protocol Semantics. In *Proceedings of USENIX Security Symposium*, August 2001.

[15] Hung-Yun Hsieh and Raghupathy Sivakumar. A transport layer approach for achieving aggregate bandwidths on multi-homed mobile hosts. In *ACM Mobicom*, September 2002.

[16] Anthony Jones and Jim Ohlund. *Network Programming for Microsoft Windows*. Microsoft Publishing, 2002.

[17] J. Klensin. *Simple Mail Transfer Protocol (RFC 2821)*, April 2001.

[18] C. Kreibich and J. Crowcroft. Honeycomb: Creating Intrusion Detection Signatures Using Honeypots. In *HotNets-II*, 2003.

[19] David Litchfield. Defeating the stack based buffer overflow prevention mechanism of microsoft windows 2003 server. http://www.nextgenss.com/papers.htm, September 2003.

[20] G. Robert Malan, David Watson, and Farnam Jahanian. Transport and application protocol scrubbing. In *Proceedings of IEEE Infocom*, 2000.

[21] P. J. McCann and S. Chandra. PacketTypes: Abstract Specification of Network Protocol Messages. In *Proceedings of ACM SIGCOMM*, 2000.

[22] David Moore, Vern Paxson, Stefan Savage, Colleen Shannon, Stuart Staniford, and Nicholas Weaver. Inside the Slammer Worm. http://www.computer.org/security/v1n4/j4wea.htm, 2003.

[23] David Moore, Colleen Shannon, and Jeffery Brown. Code-Red: a case study on the spread and victims of an Internet worm. In *ACM Internet Measurement Workshop (IMW)*, 2002.

[24] Microsoft Security Bulletin MS03-026, September 2003. http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/bulletin/MS03-026.asp.

[25] S. W. O'Malley, T. A. Proebsting, and A. B. Montz. USC: A Universal Stub Compiler. In *Proceedings of ACM SIGCOMM*, 1994.

[26] Vern Paxson. *Flex - a scanner generator - Table of Contents*. http://www.gnu.org/software/flex/manual/.

[27] Vern Paxson. Bro: A System for Detecting Network Intruders in Real-Time. In *Computer Networks*, Dec 1999.

[28] Jonathan Pincus and Brandon Baker. Mitigations for Low-level Coding Vulnerabilities: Incomparability and Limitations. http://research.microsoft.com/users/jpincus/mitigations.pdf, 2004.

[29] J. Postel and J. Reynolds. *Telnet Protocol Specification (RFC 854)*, May 1983.

[30] J. Postel and J. Reynolds. *RFC 765 - File Transfer Protocol (FTP)*, October 1985.

[31] Niels Provos. A Virtual Honeypot Framework. Technical Report CITI-03-1, Center for Information Technology Integration, University of Michigan, October 2003.

[32] Thomas H. Ptacek and Timothy N. Newsham. Insertion, evasion, and denial of service: Eluding network intrusion detection, January 1998. http://www.insecure.org/stf/secnet_ids/secnet_ids.html.

[33] Eric Rescorla. Security holes... Who cares? In *Proceedings of USENIX Security Symposium*, August 2003.

[34] *DCE 1.1: Remote Procedure Call*. http://www.opengroup.org/onlinepubs/9629399/.

[35] W32.Sasser.Worm, April 2004. http://securityresponse.symantec.com/avcenter/venc/data/w32.sasser.worm.html.

[36] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson. *RTP: A Transport Protocol for Real-Time Applications (RFC 1889)*, January 1996.

[37] Umesh Shankar and Vern Paxson. Active Mapping: Resisting NIDS Evasion Without Altering Traffic. In *Proceedings of IEEE Symposium on Security and Privacy*, May 2003.

[38] Richard Sharpe. Server message block. http://samba.anu.edu.au/cifs/docs/what-is-smb.html.

[39] Sumeet Singh, Cristian Estan, George Varghese, and Stefan Savage. The EarlyBird System for Real-time Detection of Unknown Worms. Technical Report CS2003-0761, University of California at San Diego, 2003.

[40] Microsoft security bulletin ms02-039, January 2003. http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/bulletin/MS02-039.asp.

[41] The Open Source Network Intrusion Detection System. http://www.snort.org/.

[42] Stuart Staniford, Vern Paxson, and Nicholas Weaver. How to 0wn the Internet in Your Spare Time. In *Proceedings of the 11th USENIX Security Symposium*, August 2002.

[43] Peter Szor and Peter Ferrie. Hunting for Metamorphic. Symantec Security Response.

[44] David Wagner, Jeffrey S. Foster, Eric A. Brewer, and Alexander Aiken. A First Step Towards Automated Detection of Buffer Overrun Vulnerabilities. In *NDSS*, 2000.

[45] Nicholas Weaver, Vern Paxson, Stuart Staniford, and Robert Cunningham. Large Scale Malicious Code: A Research Agenda. http://www.cs.berkeley.edu/~nweaver/large_scale_malicious_code.pdf, 2003.

[46] Nicholas Weaver, Stuart Staniford, and Vern Paxson. Very Fast Containment of Scanning Worms, 2004. http://www.icsi.berkeley.edu/ nweaver/containment/.

[47] Nick Weaver. The potential for very fast internet plagues. http://www.cs.berkeley.edu/~nweaver/warhol.html.

[48] Matthew M. Williamson. Throttling viruses: Restricting propagation to defeat malicious mobile code. Technical Report HPL-2002-172, HP Labs Bristol, 2002.

[49] Rafal Wojtczuk. Defeating Solar Designer's Non-executable Stack Patch. http://www.insecure.org/sploits/non-executable.stack.problems.html, January 1998.