

FIDRAN: A Flexible Intrusion Detection and Response Framework for Active Networks

A. Hess, M. Jung, G. Schäfer
Telecommunication Networks Group, Technische Universität Berlin
Einsteinufer 25, 10587 Berlin, Germany
Email: [hess,mjung,schaefer]@ee.tu-berlin.de

Abstract

Securing communication networks can no longer be ensured by singular and isolated security technologies like Internet firewalls or intrusion detection systems but rather calls for a combination of existing and emerging detection and response mechanisms, e.g. DDoS response mechanisms, anomaly detection, honey pots, etc. Today, most current systems prove to be too static to provide an adequate platform for a constructive teamwork of different security technologies. Therefore, we developed the FIDRAN framework for flexible intrusion detection and response that is based on an underlying active networking environment and that allows to dynamically combine existing and emerging security technologies.

FIDRAN follows a highly modular approach that allows to extend the functionality of the framework by the integration of new security modules which are implemented as active networking services, making use of next generation networks capabilities like dynamic distribution and deployment of services on active network nodes. A further advantage of the realization of FIDRAN on top of an active networking environment is the simplification and automation of maintenance work and configuration tasks.

1. Introduction

Recent developments show that securing communication networks with singular and isolated techniques proves to be insufficient to cope with the vulnerabilities of today's networks in a timely manner. The reasons behind this trend originate from multiple developments. First, the steadily increasing number of hosts connected to the Internet implying an accordingly increasing number of vulnerable hosts offers an ever growing number of potential targets for malicious activities. Second, many private and professional users are not sensible to security vulnerabilities affecting their own

machines or they are just overstrained patching these. Furthermore, many users believe that they will never become the target of an attack, due to irregular on-line times, changing IP-addresses or having the perception that their system or data, respectively, is not of value for potential hackers. Unfortunately, this is not true: As, for example, Lance Spitzner writes in his book [10]: "On February 28, 1999, at 20:15 I put the honey pot online ... Within 15 minutes of my connecting the honey pot to the Internet, an attacker identified, probed, and exploited it". Beyond this, he states that a home network was scanned on average by 31 different systems a day in the beginning of 2002.

Another reason for the rising danger arising of malicious activities is the alarming evolution of the execution speed of computer attacks. Consequently, the time window to invoke countermeasures in order to limit the harm of an attack is shrinking [4]. Weaver claims in [12] that it is possible to construct hyper-virulent active worms which are capable of infecting all vulnerable hosts of the Internet in approximately 15 minutes to an hour. Furthermore, the authors of [11] argue that under certain conditions a small worm "can infect almost all vulnerable servers on the Internet in less than thirty seconds".

As can already be seen from this short abstract of security problems in current communication networks, the currently existing security technologies on their own are not capable to react or be adapted to new attacks and changing requirements in a sufficiently timely manner. However, as each security technology has its specific advantages and drawbacks, a security infrastructure is needed that is able to combine as many technologies as possible to minimize their drawbacks and to combine their strengths.

In order to realize such a security infrastructure we developed *FIDRAN - a Flexible Intrusion Detection and Response Framework for Active Networks*. This framework allows the cooperation of traditional (firewall, intrusion detection system, etc.), innovative (DDoS defense, honey pots, etc.) and emerging security technologies in order to adequately secure communication networks. FIDRAN is build

on top of an underlying active networking environment which allows to dynamically deploy new security modules on FIDRAN-hosts. Its modular design provides the infrastructure for a constructive cooperation between modules of different security technologies. In addition, the active networking infrastructure facilitates maintenance work and enables the distribution of security tasks among different FIDRAN-hosts.

2 FIDRAN

The active networking environment which is the basis of FIDRAN consists of active nodes, a service repository (SR), a network administrator and end-systems (see figure 1). An active node is able to execute services which can dynamically be downloaded from the SR. Figure 1 illustrates an example topology consisting of three subnets, a service repository and one administrator for the three subnets. Each subnet in turn consists of several end-systems and an active node which is the gateway between subnet and the Internet. Subnet n contains a further active node on which simultaneously the local mail server is running. Each FIDRAN-node is supplied with a *security policy* which is explained in detail in section 2.1.4. The download of an op-module (see section 2.1.2) from the SR to a FIDRAN node is depicted in figure 1. Each op-module consists of the module itself, a description and a digital signature. Digital signature and description are the required input values for the initial test which each FIDRAN node performs on any op-module before integrating it.

FIDRAN is built upon the described active networking infrastructure. We posed the following design-requirements:

- subnet / node specific protection → distribution of security tasks in order to:
 - to scale the amount of security tasks to be performed per FIDRAN-host
 - to limit the amount of attack signatures per FIDRAN-host
 - to detect insider attacks
- modular concept:
 - combination of different security technologies
 - dynamic extension of the functionality through the integration of new modules
 - facilitation and acceleration of maintenance work and configuration tasks
- integration of third parties security modules
- efficiency

- (re-) configuration of the FIDRAN-system at runtime

Subnet or node specific protection, respectively, considers the idea depicted in figure 1. FIDRAN could either be running on a single node, e.g. on the gateway between a subnet and the Internet or it could be running on several hosts of a subnet. In the former case, the FIDRAN-gateway is performing all security tasks on its own, whereas in the latter case, the security operations to be performed by FIDRAN are distributed among several hosts. The specification which security operations to execute on what traffic is specified in the security policy (see section 2.1.4).

The fact that vulnerabilities and attacks can be categorized with reference to their potential victims (e.g. OS, application software) allows the realization of a demand driven intrusion detection and response infrastructure, as it usually makes little sense to scan traffic for attacks targeted at a Windows OS, while the subnet to be protected consists solely of Linux hosts. However, if a network administrator intends to scan for vulnerabilities of multiple operating systems, this is, of course, possible.

With the aid of the security policy the network administrator is able to specify and to distribute security duties among FIDRAN-hosts. Regarding subnet n in figure 1 the gateway scans the traffic which is addressed to the subnet but not to the mail-server. The mail server itself examines this traffic. Consequently, security responsibilities and the thereto appertaining work are distributed among the involved FIDRAN-hosts. This feature allows to individually scale the amount of network traffic to be analyzed in depth per FIDRAN-host which again results in a more efficient supervision. Furthermore, the FIDRAN-system becomes more resistant against evasion techniques. Among others each FIDRAN-host observes the traffic for a defined set of attack signatures and consequently they do not react to every "malicious packet" which can be easily generated in huge masses by an IDS stress tool like Stick [7].

FIDRAN consists of a management module, a control module, a security policy and a varying set of operational modules (op). An op-module adds security functionality to the FIDRAN system, whereas management and control module are responsible for configuration and administrative issues. The security policy provides the necessary information that is required by the management and control module. An op-module can be an attack signature detection module, an anomaly detection module or a response mechanism. Each op-module is designed as an active networking service, consequently each op-module can dynamically be downloaded and installed on a FIDRAN node.

Besides this, FIDRAN allows the deployment of third party's user space op-modules. Thereby, a copy of each specified packet addressed to the user's host is transferred to the third party's op-module. Hence, a user is able to develop and to deploy his own op-module on a close or remote

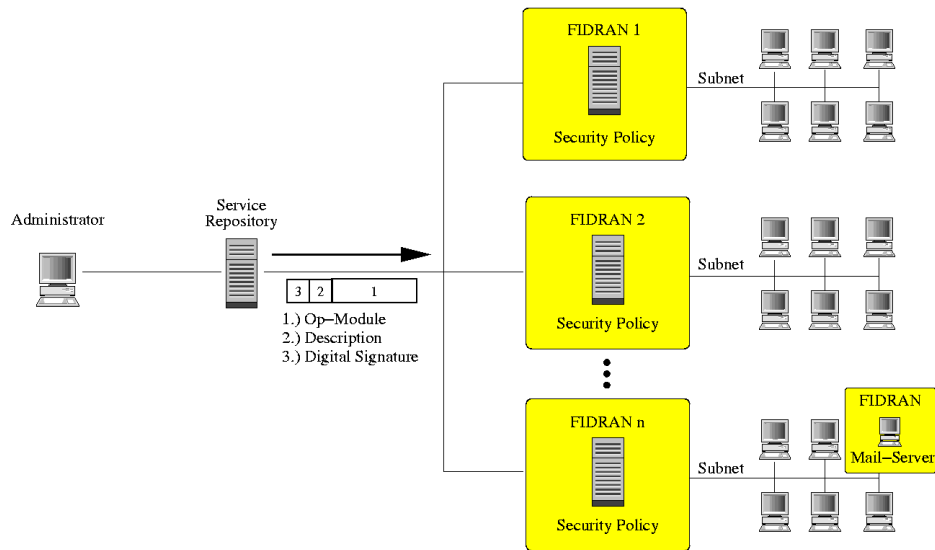


Figure 1. Distribution of Security Tasks

FIDRAN host. Such a module is authorized to send alarms to the management module and to the user's host but the initialization of countermeasures must have been previously authorized by the network administrator. The execution of third party op-modules is supervised by the integrated access control and resource monitoring mechanism which is described in detail in [9].

Further on, it is possible to configure the control module such that a copy of each alarm is sent to another specified op-module or user-space process which collects and analyzes the alarms. For example if the rate of executed alarms within a defined period of time increases in a tremendous manner, this op-module executes an alarm which again initiates a specified response mechanism.

An efficient supervision is the basis for a real-time intrusion detection and response system. Thus, highly trusted modules (e.g. according to the authorship) can be executed in kernel space whereas third party modules are executed in user space. Furthermore, the feature to distribute security operation among different FIDRAN-hosts allows to individually scale the amount of work for each host.

The reconfiguration of FIDRAN host is enabled through the integrated policy framework which is described in [8].

2.1 The FIDRAN Architecture

Figure 2 depicts the FIDRAN architecture that consists of a management module, a control module, a security policy and a varying set of FIDRAN op-modules (anomaly detection, signature detection, bandwidth monitor, etc.).

2.1.1 The FIDRAN-Management Module

The management module constitutes the interface between active networking software and the FIDRAN system. In detail the management module is responsible for the initial testing, the loading / unloading of op-modules as well as triggering response mechanisms.

An op-module which should be integrated into a FIDRAN system is transferred from the active node management module to the FIDRAN management module (see figure 2). Before the op-module can be integrated into the system, the FIDRAN management module performs an initial check on it. Thereby the management module verifies if the op-module is correctly digitally signed by the network administrator or another trusted source and if the security policy authorizes / requests the integration of the module. For a more detailed description of the initial check, please refer to [8].

After having successfully passed the initial check, the management module loads the op-module into the kernel. The FIDRAN management module is the only process which is authorized to load /unload kernel modules. This is supervised by the integrated access control mechanism [9].

Furthermore, the FIDRAN management module is responsible for the initiation of the user-space parts of the response mechanism (see section 2.1.5).

2.1.2 A FIDRAN Operational (op-) Module

Each operational module performs an individual set of operations on a packet and returns the result to the FIDRAN control module. Consequently, op-modules are realized as active networking services, such that each op-module can

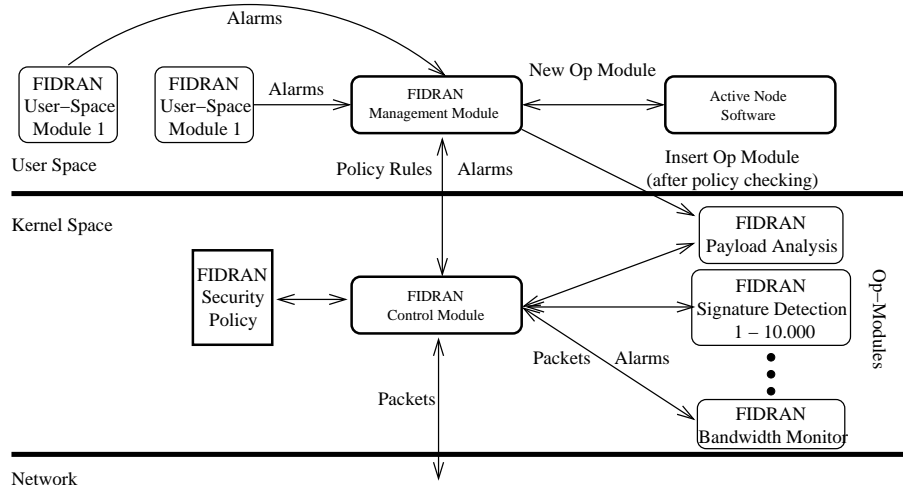


Figure 2. The FIDRAN Architecture

dynamically be downloaded and integrated into the running FIDRAN system. Thereby, the possibility to influence the execution sequence of the op-modules is given through labelling them with a priority. Thus, op-modules which scan packets for frequent and highly malicious attacks should be labelled with a high priority. The priority label is specified by the network administrator. Further on, each FIDRAN op-module contains a description which specifies its competence. A FIDRAN op-module could either be a signature detection module, an anomaly detection module or any other kind of security module.

2.1.3 The FIDRAN-Control Module

The control module constitutes the central unit of FIDRAN in kernel space and is responsible for the registration and deregistration of op-modules. Furthermore, it distributes the traffic packets to the correct op-modules or third party's module in user space and finally, it triggers the response mechanism in kernel space and forwards the alarms to the management module.

Primary, the control module manages the FIDRAN op-modules in kernel space. The FIDRAN management module in user space loads (unloads) the op-modules. Thereby, each op-module calls a specific registration (deregistration) function which passes the information required for the FIDRAN registration (deregistration) process to the control module. The control module manages the op-modules according to figure 3. It stores a varying set of linked lists according the mapping of traffic streams onto security requirements which is specified in the security policy. In figure 3 FIDRAN differs between three traffic streams Other, TCP and UDP. It is also possible to define a linked list for one specific host. Further on, the linked list are sorted accord-

ing to the priorities of the op-modules. The first element of the TCP list points to the op-module which contains the detection algorithms for the highest prioritized TCP specific attacks. In the given example of figure 3, a TCP packet would traverse all TCP op-modules until reaching the end of the list or an attack is detected. An op-module can be the member of one or more lists as certain attacks are protocol independent. Besides this, the control module is the unit that listens to the network traffic. Additionally, the control module specifies a set of predefined variables in order to minimize the amount of operations to be performed per op-module (IP-header, source address, etc.).

In the case that an op-module detects an attack it sends an alarm to the control module. The control module forwards the alarm to the management module and further on, it decides according to the security policy what to do with the packet (drop, forward, etc.).

2.1.4 The FIDRAN Security Policy

FIDRAN's security policy is realized according to the approach presented in [8]. Thereby, one particularity is the possibility to change the security policy at runtime. The security policy defines the following:

- specification of FIDRAN op-modules that can be loaded
- mapping of traffic stream \rightarrow security operations
- mapping of attack \rightarrow response mechanism
- alarm forwarding rules

Primary the FIDRAN security policy specifies the op-modules which could be integrated into the local FIDRAN

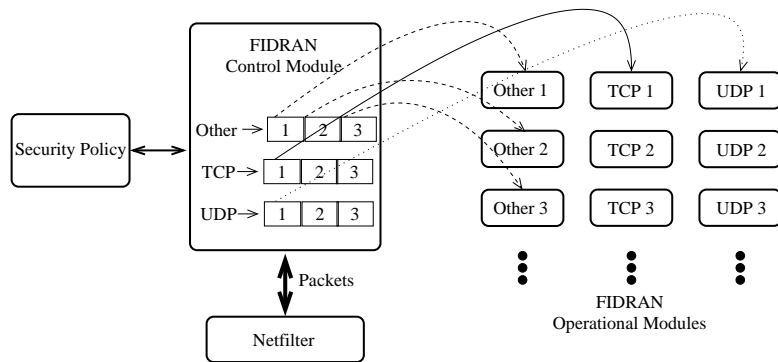


Figure 3. Packet Distribution

system. Each op-module contains a description of its competency (see figure 1) which is compared during the initial check with the local FIDRAN security policy. An op-module can only be loaded in the case that the security policy authorizes this. Besides preventing the uncontrolled introduction of potentially malicious op-modules into a FIDRAN node, this serves the purpose to keep the amount of op-modules in kernel space as small as possible. The network administrator specifies in the security policy which traffic should be observed in what manner. For instance a FIDRAN host is assigned to observe all smtp-traffic. Then the security policy of this host would only authorize the installation of smtp-specific op-modules.

Further on, a FIDRAN host could be assigned to observe many different traffic streams and therefore a mapping must be made in the policy specifying which traffic stream requires which treatment. For instance FIDRAN is running on a gateway to the Internet. Thereby, the network administrator could specify to observe any incoming traffic in a complete different manner than outgoing traffic.

Additionally, the network administrator defines how to react to an alarm, specifying the immediate reaction, what to do with packets that generate an alarm and also defining more elaborated and time consuming reactions like the initialization of a new active service, etc.

Alarm forwarding allows the evaluation of the collectivity of all local alarms. This means that the control module forwards all alarms to a specified op-module / user space process which again is capable to execute an alarm of a higher priority. For instance the module analyzes the alarms with respect to a correlation between them or it just measures the amount of alarms executed within a defined period of time.

2.1.5 Response Mechanisms

Response mechanisms to attacks are specified in the local FIDRAN security policy. Potential reactions include simply

discarding packets, firewall reconfiguration, traffic redirection, or the invocation of additional active services.

A response mechanism normally consists of a kernel space and user space part. In kernel space the control module decides according the security policy what to do with the packets that caused the alarm (drop, accept, etc.). Further on, the management module triggers further response activities as e.g. a traceback mechanism in case of a detected DDoS-attack which uses spoofed addresses.

3 Measurements

In the following we present preliminary results of a FIDRAN prototype realized on top of the active networking environment AMnet [1] running on Pentium III 800 machines with Linux 2.4.

A file of length 128 MByte was transmitted via FTP from a server to a client interconnected with a Fast Ethernet network. The FTP-transfer was monitored by a FIDRAN host located between server and client. We varied the number of installed op-modules (from 0 - 500), where each op-module was a so called null-filter, that checks for every packet if it is a TCP segment and, if so, whether at least one of the TCP-flags SYN, ACK, FIN, RST, URG, or PSH is set. Figure 4 depicts the experiment results.

We measured a relative overhead of 8% / 147% in case of 100 / 500 installed op-modules. As can already be deduced from this, the supervision can cause significant workload, depending on the amount and complexity of the installed op-modules, therefore demanding for a demand driven approach as FIDRAN.

4 Related Work

Snort [2] is able to collect network traffic and compares it with known attack signatures. However, if the attack is sufficiently distributed (spatially / temporally), then in most cases Snort will not detect the attack.

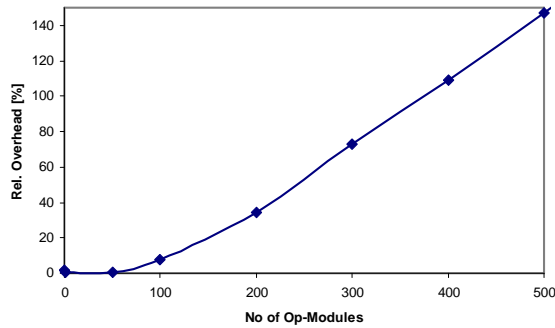


Figure 4. Relative IDS Overhead

The paper Active Network Based DDoS Defense [5] describes how active networking technology can be used for DDoS protection. The presented approach consists of a sensor which remarks a rapid increase of network traffic and a mobile traffic rate limiter which clones itself. The rate limiter migrates upstream along the attack path in order to stem the attack.

The Intrusion Blocker based on Active Networks - IBAN [6] consists of a management station, mobile vulnerabilities scanners, and mobile intrusion blockers. A mobile scanner is an application designed to detect one particular vulnerability by looking at system fingerprints. If the scanner has found a vulnerable service an intrusion blocker is placed close to the corresponding system which inspects the traffic for the vulnerable service and blocks the traffic if it detects an attack attempt. IBAN focuses on the detection of automated known attacks. A scanner and a blocker are designed for one particular vulnerability. A mobile application is designed for a particular vulnerability. Consequently, numerous mobile applications could exist in an average network. Further on, each application observes the traffic for a specific traffic pattern, thus each mobile application performs a set of identical operations

The FLAME project [3] allows users to install kernel modules for real-time packet monitoring. The code is written in Cyclone and is processed by a trusted compiler. A set of credentials is used at compile time to verify that the module is authorized to perform the requested actions. Even if the compiler could be trusted it is still dangerous to install user modules in kernel space. In contrast FIDRAN allows the loading of kernel modules which originates from trusted sources. Further on, a control module coordinates the kernel modules which improves efficiency.

Summarizing, we state that few projects exist which exploit the possibilities provided by an active networking environment or which allow the integration of different security technologies.

5 Conclusions

FIDRAN allows the integration of different security technologies in order to secure communication networks against attacks. Its modular concept provides an adequate infrastructure to dynamically add a new security function to the system. With the aid of the underlying active networking environment the administrative tasks are simplified. For instance the deployment of new FIDRAN op-modules happens in a highly automated way. The integrated policy framework allows to distribute and to scale security operations. Another feature of FIDRAN is that third parties can deploy their own modules in user space. According to their trustfulness they have certain rights to invoke countermeasures. Third party modules are supervised by an integrated access control and resource monitoring mechanism.

References

- [1] Flexinet. <http://www.flexinet.de>.
- [2] Snort. <http://www.snort.org>.
- [3] K. G. Anagnostakis, S. Ioannidis, S. Miltchev, J. Ioannidis, M. B. Greenwald, and J. M. Smith. Efficient packet monitoring for network management. In *Proceedings of IFIP/IEEE Network Operations and Management Symposium (NOMS) 2002*, April 2002.
- [4] H. K. Browne, W. A. Arbaugh, J. McHugh, and W. L. Fithen. A trend analysis of exploitations. In F. M. Titsworth, editor, *Proceedings of the 2001 IEEE Symposium on Security and Privacy*, pages 214–231, Los Alamitos, CA, May 14–16 2001. IEEE Computer Society.
- [5] D. S. et al. Active network based ddos defense. In *Proc. of Dance 2002*, 2002.
- [6] W. L. C. et al. Iban: Intrusion blocker based on active networks. In *Proc. of Dance 2002*.
- [7] C. Giovanni. Fun with packets: Designing a stick. <http://www.eurocompton.net/stick/papers/Peopledos.pdf>.
- [8] A. Hess and G. Schaefer. A flexible and dynamic access control policy framework for an active networking environment. In *Proc. of Kommunikation in Verteilten Systemen (KiVS 2003)*, pages 321–333, Leipzig, Germany, Feb. 2003.
- [9] A. Hess, M. Schoeller, G. Schaefer, M. Zitterbart, and A. Wolisz. A dynamic and flexible access control and resource monitoring mechanism for active nodes. In *Short Paper Proc. of OpenArch 2002*, pages 11–16, New York, USA, June 2002.
- [10] L. Spitzner. *Tracking Hackers*. Addison Wesley, 2002.
- [11] G. G. Staniford, S. and R. Jonkman. Flash worms: Thirty seconds to infect the internet. <http://www.silicondefense.com/flash/>.
- [12] N. C. Weaver. Warhol worms: The potential for very fast internet plagues. <http://www.cs.berkeley.edu/~nweaver/warhol.html>.