CESNET Intrusion Detection System

PAVEL VACHEK
CESNET CERTS
CESNET, Association of Legal Entities
Zikova 4, 160 00 Prague 6
THE CZECH REPUBLIC

Abstract: This paper describes a simple system for detecting hacker- and virus-induced intrusions. Its main component is the *LaBrea* program distributed under GPL licence, running on a PC server under Linux. Auxiliary programs written in-house allow automated distribution of warning e-mails to administrators of networks where the attacks originated.

Key-Words: CESNET, DSHIELD, IDS, Intrusion Detection System, LaBrea, Linux, PC

1 Introduction

A simple and useful Intrusion Detection System (IDS) has been developed in CESNET, the National Research and Educational Network of the Czech Republic. It has been operating since December 2002, informing administrators from selected networks (hereinafter, from *own networks*) about any TCP attacks originating there [1], [2]. As a result, those administrators are informed soon; they can check the potentially compromised machines, disinfect them or disconnect them from the Internet, and thus keep up the CESNET reputation. - Attacks originating outside these *own networks* are reported to *DSHIELD*, a global cooperative security system [3].

The goal of this paper is to inform other network managers, especially those of University networks with large IP address allocations, about the system usefulness and to induce them to install a similar system to make their everyday duties much easier.

2 The IDS Project

Main components of the CESNET IDS are:

- The *LaBrea* program described in the following chapter [4].
- *LaBreaBackEnd* = program which checks the output data generated by the *LaBrea* program. Should any important records be found, *LaBreaBackEnd* searches for contact addresses of appropriate network and domain administrators, selects the preferred ones and generates files containing information about attacks.
- *LaBreaReport* = program which sends out e-mail warnings containing data generated by *LaBreaBackEnd*.
- *Labrea.pl* reads and parses output data generated by the *LaBrea* program to send them by e-mail to *DSHIELD*.

The IDS helper tools are:

- Filter for *LaBrea*-generated data using *grep*, *tee* and configuration files **LBstrings** and **LBstrings.all**;
- script **/etc/init.d/labrea** to start, restart and stop the *LaBrea* program;
- script **DShield.sh** to restart *LaBrea* and process the *LaBrea*-generated data for *DSHIELD* at regular intervals;
- script **LBbe.sh** to start the *LaBreaBackEnd* and possibly also *LaBreaReport* programs, maintain log files, etc., at regular intervals.

All helper tools are also described in detail below.

3 The *LaBrea* Program

Goal of the *LaBrea* program is to stop or at least slow down propagation of network viruses and hacker activities. Its author wrote it in response to the Code Red virus which started to propagate in July 2001. Installing just the *LaBrea* server in standalone mode helps the Internet enough; however, the complete IDS described here is much more useful.

Main features of the *LaBrea* program are:

- *LaBrea* monitors single addresses or preferably blocks of IP addresses never before assigned to end users. Therefore, one can assume safely that only hackers or network viruses attempt to connect there.
- *LaBrea* pretends that real machines exist on these monitored addresses: it creates virtual servers which accept incoming TCP connections and hold them as long as

possible so that they cannot cause harm elsewhere. These virtual servers can also respond to ICMP Echo and SYN+ACKs received.

LaBrea in the CESNET IDS responds to TCP connection attempts as follows:

- CONNECTION TRAPPING: After the TCP connection has been established successfully (SYN - SYN+ACK - ACK), *LaBrea* keeps advertising a very small TCP Receive Window. As a result, this connection never terminates by itself while requiring a very small bandwidth.
- DDoS HANDLING: Distributed Denial of Service attack packets typically contain forged source addresses. After the target machine receives a SYN packet whose source address is a part of IP space monitored by *LaBrea*, it responds with a SYN+ACK packet which arrives at the IDS. *LaBrea* responds by sending an RST packet which terminates the connection instantly and thus decreases the load of the target machine.

4 Installation

Let us suppose that our Institution has been allocated an address range 172.16.0.0/16, i.e., 65536 IP addresses. Let us also suppose that subnet 172.16.224.0/22, i.e., 1024 IP addresses, has not yet been assigned and will never be. This subnet will be reserved for our IDS interface Eth1.

The IDS interface Eth0 is assigned an address from another part of our network, e.g., 172.16.1.2. This is the only IP address advertised in DNS which needs a reverse (PTR) record. It is used for sending e-mail to administrators of *own networks* whose machines attempted to connect to the subnet monitored by the Eth1, and to send data on attacks from worldwide Internet to the *DSHIELD* project. The Eth0 interface can also be accessed using *ssh* from some selected addresses.

Before the *LaBrea* installation starts, one should check if the Linux distribution used includes programs *Gcc, Perl, Libpcap, and Whois*. The latter program is a part of SuSE Linux and several other distributions [5]. It can be downloaded from http://ftp.debian.org/debian/pool/main/w/whois. Current version is **whois_4.7.23.tar.gz** (60 kB; Sep 12, 2007).

Latest version of program *Libdnet* - probably **libdnet-1.11.tar.gz** (446 kB; Jan 19, 2006) - can be found at http://prdownloads.sourceforge.net/libdnet.

Latest version of program *LaBrea* - probably **labrea-2.5-stable-1.tar.gz** (210 kB; Oct 29, 2003) - should be downloaded from http://prdownloads.sourceforge.net/labrea.

Libdnet and *LaBrea* programs should be installed using the following command sequence:

$ **./configure**
$ **make**
$ **su**
$ **make install**

First, *LaBrea* should be run manually in Test mode (only Eth0 interface will be used, network traffic will not be influenced) using the command

$ **labrea -T -z**

In the next step, IDS should use both Ethernet interfaces:

- management interface Eth0 is unchanged
- monitored subnet 172.16.224.0/22 on Eth1, IP address can be, e.g., 172.16.225.77
- maximum bandwidth allowed = 2400 kbps
- subnet 172.16.224.0/22 can be accessed only by *LaBrea*; therefore, Auto Hard Capture and No ARP Sweep modes are selected.

Again, *LaBrea* is started manually:

$ **labrea -b -H -i eth1 -I 172.16.225.77 --no-arp-sweep -o -p 2400 -v -z**

LaBrea should start responding to TCP connection attempts, ICMP Echo (Ping) and SYN+ACKs received.

If everything works, Eth1 can be set up as unnumbered interface which further improves the IDS server security. The following paragraph which describes this procedure will help especially users of the SuSE Linux distribution; users of other distributions must use similar functions available in their operating systems.

4.1 Setting up unnumbered Eth1 interface

Original status of unconfigured Eth1 interface may be something like:

$ **ifconfig eth1**
eth1 Link encap:Ethernet HWaddr 00:09:87:65:43:21
BROADCAST MULTICAST MTU:1500 Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0 collisions:0 txqueuelen:1000
RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)

Using the *YaST* tool, appropriate interface parameters should be set up. The IP address can be

172.16.225.77 again. Afterwards, *ifconfig* may report something like:

```
$ ifconfig eth1
eth1      Link encap:Ethernet  HWaddr
00:09:87:65:43:21
          inetaddr:172.16.225.77
Bcast:172.16.227.255  Mask:255.255.252.0
          UP BROADCAST RUNNING
MULTICAST  MTU:1500  Metric:1
          RX packets:11 errors:0 dropped:0 overruns:0
frame:0
          TX packets:6 errors:0 dropped:0 overruns:0
carrier:0 collisions:0 txqueuelen:1000
          RX bytes:660 (660.0 b)  TX bytes:460
(460.0 b)
```

Configuration file for this interface located in the **/etc/sysconfig/network** directory will contain the following:

```
$ more ifcfg-eth-id-00\:09:87\:65\:43\:21
BOOTPROTO='static'
BROADCAST='172.16.227.255'
IPADDR='172.16.225.77'
MTU=''
NETMASK='255.255.252.0'
NETWORK='172.16.224.0'
REMOTE_IPADDR=''
STARTMODE='onboot'
UNIQUE='Hkyc.MvtRJ6g27Q1'
_nm_name='bus-pci-0000:01:0a.0'
```

Note: To make a backup copy of this file, one must change the start of its filename - e.g., to **sav-ifcfg-eth-id-00\:09\:87\:65\:43\:21**.

The BROADCAST, IPADDR, NETMASK and NETWORK data should be cleared using a simple text editor. Afterwards, this file will contain the following:

```
$ more ifcfg-eth-id-00\:09:87\:65\:43\:21
BOOTPROTO='static'
BROADCAST=''
IPADDR=''
MTU=''
NETMASK=''
NETWORK=''
REMOTE_IPADDR=''
STARTMODE='onboot'
UNIQUE='Hkyc.MvtRJ6g27Q1'
_nm_name='bus-pci-0000:01:0a.0'
```

Operating system or the *network* must be restarted. The following data should appear:

```
Setting up network interfaces:
   lo
   lo       IP address: 127.0.0.1/8  . . . . . done
Waiting for mandatory devices:
 eth-id-00:12:34:56:78:9a eth-id-00:09:87:65:43:21
   eth0     device: AAAAA Corp. [Ethernet ABCD]
(rev 08)
   eth0     configuration: eth-id-00:12:34:56:78:9a
   eth0     IP address: 172.16.1.2/25
   eth1     device: BBBBB Corp. [Ethernet EFGH]
(rev 09)
   eth1     configuration: eth-id-00:09:87:65:43:21
Setting up service network . . . . . done
```

One can see that the Eth1 interface is functional but no IP address is assigned. Program *ifconfig* will display its status as follows:

```
$ ifconfig eth1
eth1      Link encap:Ethernet  HWaddr
00:09:87:65:43:21
          UP BROADCAST RUNNING
MULTICAST  MTU:1500  Metric:1
          RX packets:442 errors:0 dropped:0
overruns:0 frame:0
          TX packets:10 errors:0 dropped:0
overruns:0 carrier:0 collisions:0 txqueuelen:1000
          RX bytes:3199 (3.1 Kb)  TX bytes:1005
(1.0 Kb)
```

Program *LaBrea* can be launched now using the following command line (for detailed information on command-line options please see the **man labrea** output):

```
$ labrea -b -E 00:09:87:65:43:21 -H -i eth1 -I
172.16.225.77 -n 172.16.224.0/22 --no-arp-sweep
-O -o -p 2400 -v -z
```

Program *LaBrea* should run now for some time to allow administrators to become acquainted with its features (command-line and/or configuration file parameters). One can try to connect to the monitored range using, e.g., the *Telnet* program, test how *LaBrea* responds to the *Ping* command, learn about the bandwidth consumed, frequency of attacks, number of DDoS attacks seeming to originate from the monitored subnet, etc.

4.2 Installing the remaining IDS components

One will find out that *LaBrea* generates too much unnecessary data, and most of it is coming out of foreign networks which are not too important for us. Therefore, the *LaBrea* output data will be processed in two ways: data on attacks from our *own networks* will be processed directly and often by our IDS while the remaining data will be passed over to *DSHIELD* for further processing according to its own rules.

The data to be processed directly should contain only the following:

- data on attacks originating from *own networks*. This filtering need not be perfect - exact filtering is configured using @ownnet in *LaBreaBackEnd*;
- logged packets SYN+ACK which signify DDoS attacks coming ostensibly from our *own networks*;
- logged packets ICMP Echo - not processed but useful for IDS operation monitoring;
- bandwidth reported - useful for determining the IDS usefulness. Appropriate configuration data can be found in file **LBstrings**.

Data on all connection attempts from the worldwide Internet is passed over to *DSHIELD*. Appropriate configuration data can be found in file **LBstrings.all**.

One must also make sure that

- *LaBrea* starts automatically after the operating system is loaded;
- results of *LaBrea* run are processed in regular intervals;
- the IDS can handle a system shutdown without data loss.

This is the purpose of file **/etc/init.d/labrea** which should be installed using the command

$ **insserv -v labrea**

- Command **labrea start** launches the programs *LaBrea*, *grep* and *tee*; filtered output data coming from *own networks* is appended to the existing file **LaBrea.new**. Similarly, filtered data coming from the worldwide Internet is appended to the existing file **LaBrea.all**;
- command **labrea stop** terminates the *LaBrea* program;
- command **labrea restart** given at times specified in the *crontab* file performs these functions:
    - terminates the *LaBrea* program
    - output file **LaBrea.new** is appended to **/var/log/LaBrea** and deleted
    - output file **LaBrea.all** is appended to **/var/log/LaBrea.ALL** and deleted
    - programs *LaBrea*, *grep* and *tee* are launched again; filtered *LaBrea* output data is written into new files **LaBrea.new** and **LaBrea.all**
- command **labrea status** displays the *LaBrea* status as well as the number of lines of **LaBrea.new** and **LaBrea.all** output files.

CESNET IDS processes the *LaBrea* output data at times specified by the *crontab* file of user **root**. The following operations are performed:

- script **DShield.sh** (running every second hour every day) restarts *LaBrea* and processes, sends and deletes the *DSHIELD* data file **LaBrea.ALL**
- six times every working day, script **LBbe.sh** performs these functions:
- launches program *LaBreaBackEnd* which processes the *LaBrea* output data;
- should any attacks from *own networks* be detected, program *LaBreaReport* is launched to send warning e-mails to appropriate network/domain administrators;
- the **/var/log/LaBrea** output log is renamed and compressed.

Programs *LaBreaBackEnd* and *LaBreaReport* are written in *Perl*. Both of them require module *Mail::Sender* from http://search.cpan.org/dist/Mail-Sender to be installed [6]. Currently, its latest version is **Mail-Sender-0.8.13.tar.gz** dated Feb 25, 2006. It can be installed using the following commands:

$ **perl Makefile.PL**
$ **make**
$ **make test**
$ **make install**

Several default values must be set up during *Mail::Sender* installation. The most important is the mail server address. If *Postfix* runs on the IDS server, this address is given as **localhost**. All other questions can be answered by pressing **Enter** or (in case of default coding) by pressing **n**.

The *LaBreaBackEnd* in the CESNET IDS is launched as follows:

$ **LaBreaBackEnd -b -i file1 -l limit -O -v**

-b batch mode
-i *file1* *file1* = input filename (*LaBrea* output log)
-l *limit* attacks below *limit* from a single IP address are ignored
limit = 0 ... all attacks are processed
-O only attacks from networks listed in @ownnet are processed
-v verbose output.

One purpose of *LaBreaBackEnd* is to search for e-mail addresses of all administrators responsible for networks where the attacks have originated. *LaBreaBackEnd* also finds all administrators responsible for the originating domain if a reverse record for the IP address in question is found. Program selects the most appropriate address(es) from the resulting list. Preferred addresses are generated by *LaBreaBackEnd* and added to the list

if none is found in it.

LaBreaBackEnd minimises the number of warning e-mails sent: Should *LaBrea* detect attacks from different networks managed by a single group of administrators, a single e-mail with all appropriate error reports is sent to this group.

LaBreaBackEnd generates log file **yymmdd.log** with detailed data on program run and results. Should *LaBrea* detect any attacks from *own networks*, the following is also created:

- directory **result** containing data files on detected attacks to be used by *LaBreaReport*
- archive file containing the above attack data - **yymmddhh.zip**.

LaBreaBackEnd also sends the system administrator a summary SMS text message containing the number of format errors in input data, number of attacking machines and number of generated data files.

LaBreaReport is launched only if *LaBreaBackEnd* has generated any data files in the **result** directory. *LaBreaReport* reads these files, formats them properly, adds some explanatory notes and sends them by e-mail. Implicitly, they are deleted afterwards. Various command-line options can be used; the following are used by the CESNET IDS:

$ **LaBreaReport -k -i**

-k	files are not deleted but renamed to ***.sent**
-i	files ***.sent** in **result** directory are ignored.

The following command-line parameters can additionally be used for testing *LaBreaReport*:

-t	no e-mails are sent, no files are deleted
-T	e-mails are sent only to $testaddr
-v	verbose mode.

Program *LaBreaBackEnd* contains some variables which should be checked and/or set up. These are especially the following:

@ignoredom	no warning e-mail is sent to admins of domains listed in @ignoredom
@ndom	attacks from domains listed in @ndom are ignored
@nipadd	attacks from single IP addresses listed in @nipadd are ignored
@nipnet	attacks from networks listed in @nipnet are ignored
@ownnet	*own networks* listed in @ownnet will be monitored only if option **-O** is selected
$smsadd	e-mail address where the summary SMS message is sent.

Exact IP addresses of the monitored network should not be displayed in the warning messages. The following two variables will help:

$strin1, $strin2	strings (parts of destination IP addresses) to be replaced
$strout1, $strout2	... by these strings.

The following *LaBreaReport* variables should also be set up:

$bccaddr	blind copies (Bcc:) are sent here unless the **-t** or **-T** command-line options are selected
$testaddr	address where warning e-mails are sent if option **-T** is selected.

Source files of the Perl programs as well as configuration files can be found at address ftp://ftp.cesnet.cz/local/ids signed by the author's PGP key 1BEDF25D (fingerprint = 6627 E8B9 29EC A28A EBE3 8A7F 44C4 BABC 1BED F25D).

5 Recommended IDS configuration

The following long-time experience may help in choosing a proper IDS server hardware:

- IDS does not need much CPU power.
- Bandwidth of monitored data is adjustable. Average bandwidth detected by the CESNET IDS is usually about 150 kbps; maximum bandwidth is set to 2400 kbps.
- The strongest attack detected so far between one Friday evening and Monday morning in 2005 had generated a log file **LaBrea.new** of some 130 MB (4.5 MB after bzip2 compression).
- Uncompressed data sent by e-mail to *DSHIELD* every 2 hours is usually below 10 MB; CESNET IDS does not archive it.
- One can see that extremely large disc capacities are not required, either.

Currently, the CESNET IDS runs on a DELL PowerEdge 1425SC in the following configuration:

- CPU Intel Xeon, 3400 MHz
- 1 GB RAM
- SCSI adapter Adaptec AIC7902 Ultra320
- SCSI disc FUJITSU MAT3073NP (73 GB, 10k RPM)
- Two Intel Gigabit Ethernets (one for *ssh* management access and mail, second for attack monitoring)
- SuSE Linux version 10.x operating system.

The IDS should be located as close to the Institution border router as possible. Complete TCP traffic should enter the IDS - no TCP ports should be

filtered. If a large subnet is to be monitored, no switch should be used to connect the IDS because its ARP cache might overflow.

6 Experience Gained

Originally, the CESNET IDS was designed to detect and report attacks from the worldwide Internet. As a result, many of its features are not currently used but its author wishes to keep them for those who may find them useful.

The CESNET IDS is concentrating on attacks from its *own networks* for administrative reasons only. As soon as the IDS started sending out its warning to the worldwide Internet in 2002, most of their recipients (several thousand per week) kept requesting special setup exceptions for their networks. The CESNET IDS was redesigned to cope with their requirements but still, the IDS administrator would have to spend several hours each day communicating with the administrators and updating the *LaBreaBackEnd* configuration according to their wishes.

Another big problem are the missing or obsolete administrator e-mail addresses as well as abuse addresses according to RFC 2142 (abuse@example.net). Much of the IDS administrator's work is futile if the network/domain administrator contact addresses are not updated regularly in the ARIN, APNIC, RIPE DB, etc.

This is why the author of CESNET IDS decided to concentrate on distributing the warning e-mails only to the CESNET network where he can persuade the network managers to keep their contact data in the RIPE DB up-to-date and to set up the abuse addresses, too. The remaining data on attacks from the worldwide Internet is sent by e-mail to the *DSHIELD* project.

6.1 CESNET IDS Advantages

- Excellent for institutions with large IP address assignments;
- Fast and reliable detection of virus-infected or hacked machines on "near" IP addresses;
- Detection capability can be improved by increasing size of monitored network;
- No false positives;
- Simple to install, no maintenance necessary.

6.2 CESNET IDS Disadvantages

- Rather weak detection capability for virus-infected `distant' networks (viruses prefer infecting `near' machines within their current /8 network);
- Unsuitable for UDP attack detection (of course).

7 Conclusion

The CESNET IDS described above is used for monitoring a single /16 address allocation. This is also the basic configuration suitable for most large institutions. However, CESNET is a NREN with several independent large allocations, all of which should be monitored. Therefore, for several months the CESNET IDS has been running in an advanced mode where two different large allocations are being monitored by a single IDS server. To achieve this, two independent instances of *LaBrea* are running and only simple configuration changes are necessary. The two Ethernet interfaces described above are sufficient for this task.

The IDS runs perfectly. Human intervention is necessary just when some abuse contacts become invalid or when network administrators ignore their received mail.

This author will be glad to receive a message about a newly installed IDS.

References:

[1] Pavel Vachek, Zabezpečení lokálních sítí CESNET2, Vysokorychlostní síť národního výzkumu a její nové aplikace, 2003, pp. 228-231. http://www.cesnet.cz/doc/2003/zprava/security.html

[2] Pavel Vachek, CESNET Intrusion Detection System. CESNET Technical Report 5/2006. http://www.cesnet.cz/doc/techzpravy/2006/ids/

[3] DSHIELD, How to submit your firewall logs to DShield. http://www.dshield.org/howto.html

[4] Tom Liston, LaBrea: "Sticky" Honeypot and IDS. http://labrea.sourceforge.net

[5] Marco d'Itri, Whois - an improved whois client. http://www.linux.it/~md/software

[6] Jan Krynicky, Mail::Sender - module for sending mails with attachments through an SMTP server. http://search.cpan.org/dist/Mail-Sender