Malware Analysis Report

New C&C Protocol for ZeroAccess/Sirefef
June 2012

Kevin McNamee
www.kindsight.net

Analysis Summary

Name: Trojan:Win32/Sirefef.P
MD5: c71d6136d7549559ebddf65a48dd6a06
Size: 156672 bytes
Source: CleanMX
File Type: PE32 executable for MS Windows (GUI) Intel 80386 32-bit
Sample Collected: 2012-05-28 11:18:20

Introduction

We have been investigating the appearance of a new variation of the ZeroAccess/Sirefef bot. In February, we published a detailed analysis of the network behavior of this bot and the encrypted p2p protocol that it uses to communicate with its peers. The main purpose of this botnet is to distribute malware responsible for ad-click fraud.

The traffic generated by the ad-click fraud is 0.1MBits/second when averaged out. For the infected consumer, this adds up to 32GBytes per month which it is the equivalent of downloading 45 full length movies. For the service provider, the impact on their network depends on the number of infected subscribers. The observed infection rate in mid-June was about 0.8%. This means that at any instant this bot alone is consuming 800MBits/sec of bandwidth for every 1M users on the network.

The underlying structure and function of the bot remain the same, but the command and control (C&C) protocol has switched to a combination of TCP and UDP. It also attempts to use UDP broadcasts addresses as part of its rallying strategy. We did not see any successful responses from this. The botnet continues to be very prolific with this new variety infecting about 0.8% of the home networks protected by Kindsight. Over a one week period on one network, we observed 2856 infected computers actively communicating with over one million Internet peers.

Detailed Analysis

How the New C&C Protocol Works

The bot still uses a peer-to-peer strategy for communication, but now uses a combination of UDP and TCP protocols. A file called "@" contains a list of the current peer IP addresses. As with the previous version, it contains 256 IP addresses. Every second, the infected process sends a UDP packet to one of these peers. The sample we observed in the lab used port 16464, but we have also seen ports 16465, 16470 and 16471 used in the wild.

The packet contains 16 bytes of data and is "obfuscated" using a simple XOR scheme rather than the RC4 encryption that was used previously.

CRC	Command	0	32 bit value

The CRC field is a check sum that is verified by the receiver. The "Command" field is set to "getL" as with the older version of the protocol. The malware command parser also looks for "retL" and "netL" as commands. The "32-bit-value" appears to identify the bot instance sending the request. The XOR scheme used to obfuscate the data involves initializing a register with the text "ftp2" and then looping through the data 4 bytes at a time XORing it with the value in the register. The register is then rotated one bit to the left on each pass through the loop. This is not exactly leading-edge cryptography.

The peer responds with the following packet that contains updates to the list of peers and a list of the additional malware files the peer has installed. This data is also obfuscated using the same XOR encoding scheme.

CRC	Command	Length	32 bit value
N1			
IP(1)	T(1)		
…	…		
IP(n1)	T(n1)		
N2			
Name1	Date1	Size1	
32 word signature			
Name2	Date2	Size2	
32 word signature			
Name3	Date3	Size3	
32 word signature			

The "Command" in the header is always "retL".

N1 is the number of IP addresses returned. Rather than send them all at once, most peers responded with only 16. These are used to update the "@" file and maintain a current list of active peers. An additional 32 bit is provided with each IP address (T(n)). In the previous version this was the time in seconds since the IP entry was updated, however in this case they were set to zero in the traffic we decoded.

The IP address list contained a number of broadcast addresses. Technically UDP supports IP broadcasts and a bot deployed on an addresses subnet could respond to the broadcast UDP packet. However, in reality this is extremely unlikely and we saw no responses to these broadcasts.

N2 contains the number of files the peer has available. In observed exchanges, this ranged from 1 to 3. The file name is stored as a 32 bit binary value and used as an 8 character hex string. File names observed included "800000cb", "00000001" and "80000000". These are similar to the names used in the previous version of the malware. These files contain the additional malware that is used for ad-click fraud and other purposes. The name is followed by what appears to be the creation date of the file and its size in bytes. Then there are 32 words of binary data that seems to be associated with the file. This may be some sort of signature that can be used to verify the correct file is uploaded.

The infected host checks the file name and creation date against what it already has. If it does not have an up-to-date copy of the file, it initiates a TCP session to the peer on the same port number as the UDP exchange and downloads the file. The file is then dropped into the U directory on the infected machine.

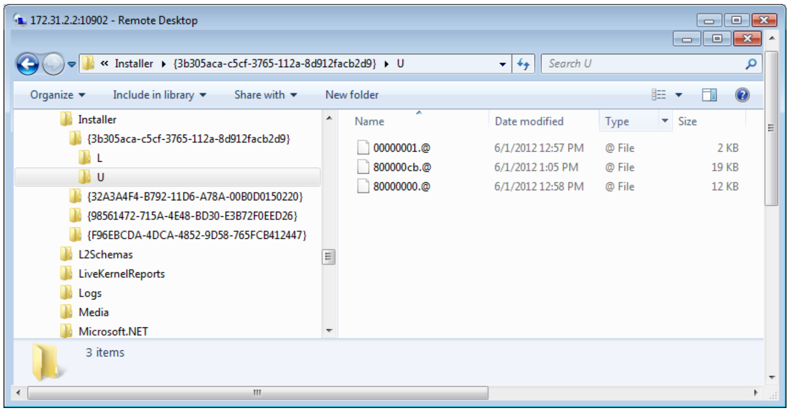
172.31.2.2:10902 - Remote Desktop
« Installer ▸ {3b305aca-c5cf-3765-112a-8d912facb2d9} ▸ U
Search U
Organize
Include in library
Share with
New folder
3 items

Installer
{3b305aca-c5cf-3765-112a-8d912facb2d9}
L
U
{32A3A4F4-B792-11D6-A78A-00B0D0150220}
{98561472-715A-4E48-BD30-E3B72F0EED26}
{F96EBCDA-4DCA-4852-9D58-765FCB412447}
L2Schemas
LiveKernelReports
Logs
Media
Microsoft.NET

Name	Date modified	Type	Size
00000001.@	6/1/2012 12:57 PM	@ File	2 KB
800000cb.@	6/1/2012 1:05 PM	@ File	19 KB
80000000.@	6/1/2012 12:58 PM	@ File	12 KB

An Updated Infection Process

The infection process and the mechanism for remaining persistent between boots of the infected computers have changed somewhat from the previous version. As before, the malware relies on social engineering, drive by downloads and exploit kits for distribution. When the malware is executed, it sets itself up in two directories using a device specific CLSID as the directory name. For example:

```
C:\Windows\Installer\{3b305aca-c5cf-3765-112a-8d912facb2d9}
C:\Users\UserName\AppData\Local\{3b305aca-c5cf-3765-112a-8d912facb2d9}
```

It drops a copy of two files into these directories.

n	- the malware executable
@	- the list of peer IP addresses

Subdirectories U and L are also created. U will contain additional downloaded malware. It is not clear what L is for.

The malware ensures it is reloaded on boot by modifying the CLSID entries in the registry for wbemess.dll and shell32.dll to load the malware instead of the normal DLLs.

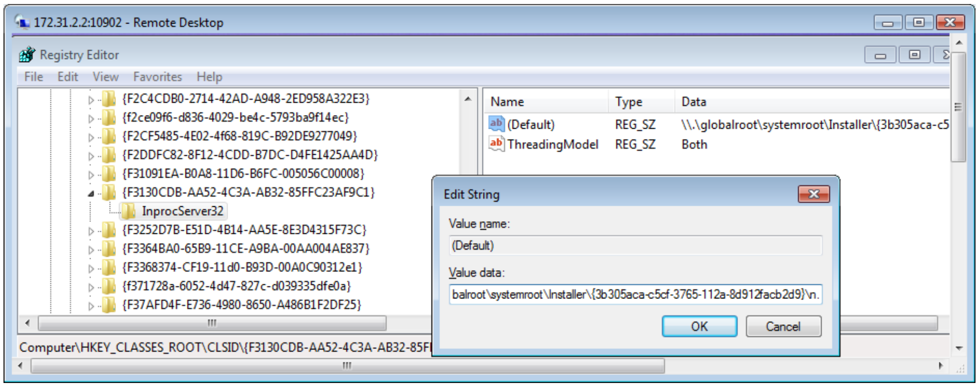

```
{F2C4CDB0-2714-42AD-A948-2ED958A322E3}
{f2ce09f6-d836-4029-be4c-5793ba9f14ec}
{F2CF5485-4E02-4f68-819C-B92DE9277049}
{F2DDFC82-8F12-4CDD-B7DC-D4FE1425AA4D}
{F31091EA-B0A8-11D6-B6FC-005056C00008}
{F3130CDB-AA52-4C3A-AB32-85FFC23AF9C1}
    InprocServer32
{F3252D7B-E51D-4B14-AA5E-8E3D4315F73C}
{F3364BA0-65B9-11CE-A9BA-00AA004AE837}
{F3368374-CF19-11d0-B93D-00A0C90312e1}
{f371728a-6052-4d47-827c-d039335dfe0a}
{F37AFD4F-E736-4980-8650-A486B1F2DF25}
```

Name	Type	Data
(Default)	REG_SZ	\\.\globalroot\systemroot\Installer\{3b305aca-c5
ThreadingModel	REG_SZ	Both

Value name: (Default)

Value data: balroot\systemroot\Installer\{3b305aca-c5cf-3765-112a-8d912facb2d9}\n.

Computer\HKEY_CLASSES_ROOT\CLSID\{F3130CDB-AA52-4C3A-AB32-85F

The malware in the \Installer directory is associated with the wbemess.ddl registry entry. It attaches itself to an svchosts process and is active in the p2p communication. We can see that process listening on the UPD and TCP ports, updating the “@” file and downloading additional malware. The malware in the \AppData directory is associated with shell32.dll registry entry and attaches itself to the explorer process (as did the original version of ZeroAccess), but remains inactive. It does not participate in the p2p protocol, does not listen on any ports and does not download any additional files. If the original process is killed or disabled this process becomes active and takes over operation of the bot.

Network Behavior During Infection Process

There is also some interesting network traffic during the infection process. The malware contacts /geo/txt/city.php at promos.fling.com to find out its location. It used 8.8.8.8 to resolve the host name. It then makes a number of connections to /count.php at livecounter.co, passing in a unique id. Interestingly, this same id used is also passed to the server where the malware picks up its ad-click instructions. There are also some malformed DNS lookup requests sent to 66.85.130.234. The 20 bytes of data looks encrypted, but elicited no response from the server. The malware sample tested also attempted to download and install an update to the Adobe Flash player, presumably in an attempt to gain escalated privileges. Once the infection is complete the traffic is mostly due to the ad-click malware and the UDP C&C.

Making Money through Ad-click Fraud

The malware wakes up periodically and browses web sites in what presumably is an attempt to make money through an ad-click scheme. The details of this activity are also a bit different from the original malware that was analyzed in February. The process is initiated by a connection to a predefined server to pick up instructions on where to browse. There was no DNS lookup for this site, so the IP address is hardcoded in the malware itself.

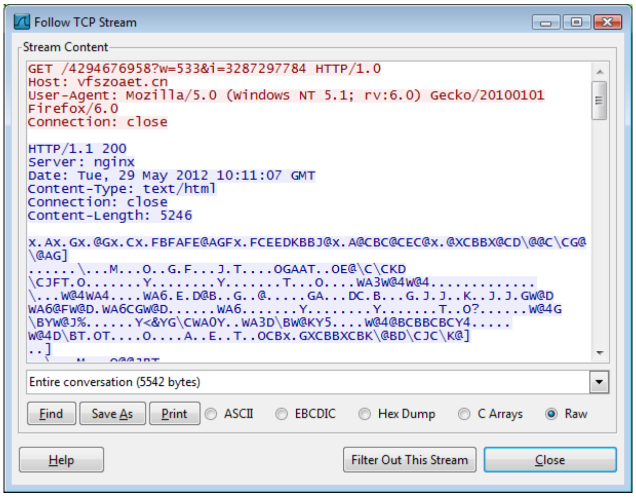

```
GET /4294676958?w=533&i=3287297784 HTTP/1.0
Host: vfszoaet.cn
User-Agent: Mozilla/5.0 (Windows NT 5.1; rv:6.0) Gecko/20100101
Firefox/6.0
Connection: close

HTTP/1.1 200
Server: nginx
Date: Tue, 29 May 2012 10:11:07 GMT
Content-Type: text/html
Connection: close
Content-Length: 5246

x.Ax.Gx.@Gx.Cx.FBFAFE@AGFx.FCEEDKBBJ@x.A@CBC@CEC@x.@XCBBX@CD\@@C\CG@
\@AG]
......\...M...O..G.F...J.T....OGAAT..OE@\C\CKD
\CJFT.O.......Y........Y.......T...O....WA3W@4W@4.............
\...W@4WA4....WA6.E.D@B..G..@.....GA...DC.B...G.J.J..K..J.J.GW@D
WA6@FW@D.WA6CGW@D......WA6.......Y........Y.......T..O?......W@4G
\BYW@J%......Y<&YG\CWA0Y..WA3D\BW@KY5....W@4@BCBBCBCY4.....
W@4D\BT.OT....O....A..E..T..OCBx.GXCBBXCBK\@BD\CJC\K@]
..]
```

Entire conversation (5542 bytes)

○ ASCII ○ EBCDIC ○ Hex Dump ○ C Arrays ● Raw

The number provided in the “w” parameter is the same as was used in the “id” parameter when contacting to livecounter.co, so the two must be related. The 5K of data returned is certainly obfuscated. The binary values are mostly in the range 0x00-0x5F, with a few exceptions.

The malware then makes a connection to another computer on TCP port 12758 and sends it a copy of the data retrieved from the exchange above. This second system responds with additional data (50K), similarly encoded. This appears to be the instructions (URLs probably) for the ad-click scheme. The infected computer then proceeds to visit hundreds of web sites during the next minute clicking on links. It then stops abruptly, waits for 5 minutes and then repeats the process from the start.

The browsing only consumes about 0.1 MBits/second when averaged over a long period. However for an individual user this adds up to 32GigBytes per month, which can have a significant impact for users with a bandwidth cap. To put it in perspective, it is the equivalent of downloading 45 full length movies. For the service provider, the impact on their network depends on the infection rate. The observed infection rate in mid June was about 0.8%. This means that at any instant this bot alone is consuming 800Mbits/sec of bandwidth for every 1M users in the network.