

Internet Engineering Task Force (IETF)
Request for Comments: 6561
Category: Informational
ISSN: 2070-1721

J. Livingood
N. Mody
M. O'Reirdan
Comcast
March 2012

Recommendations for the Remediation of Bots in ISP Networks

Abstract

This document contains recommendations on how Internet Service Providers can use various remediation techniques to manage the effects of malicious bot infestations on computers used by their subscribers. Internet users with infected computers are exposed to risks such as loss of personal data and increased susceptibility to online fraud. Such computers can also become inadvertent participants in or components of an online crime network, spam network, and/or phishing network as well as be used as a part of a distributed denial-of-service attack. Mitigating the effects of and remediating the installations of malicious bots will make it more difficult for botnets to operate and could reduce the level of online crime on the Internet in general and/or on a particular Internet Service Provider's network.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6561>.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. Key Terminology	3
1.1.1. Malicious Bots, or Bots	3
1.1.2. Bot Networks, or Botnets	4
1.1.3. Host	5
1.1.4. Malware	5
1.1.5. Fast Flux	5
2. Problem Statement	6
3. Important Notice of Limitations and Scope	7
4. Detection of Bots	8
5. Notification to Internet Users	12
5.1. Email Notification	13
5.2. Telephone Call Notification	13
5.3. Postal Mail Notification	14
5.4. Walled Garden Notification	14
5.5. Instant Message Notification	16
5.6. Short Message Service (SMS) Notification	16
5.7. Web Browser Notification	17
5.8. Considerations for Notification to Public Network Locations	18
5.9. Considerations for Notification to Network Locations Using a Shared IP Address	18
5.10. Notification and End User Expertise	19
6. Remediation of Hosts Infected with a Bot	19
6.1. Guided Remediation Process	21
6.2. Professionally Assisted Remediation Process	22
7. Failure or Refusal to Remediate	23
8. Sharing of Data from the User to the ISP	23
9. Security Considerations	23
10. Privacy Considerations	24
11. Acknowledgements	24
12. Informative References	26
Appendix A. Examples of Third-Party Malware Lists	28

1. Introduction

This document contains recommendations on how Internet Service Providers can use various remediation techniques to manage the effects of malicious bot infestations on computers used by their subscribers. Internet users with infected computers are exposed to risks such as loss of personal data and increased susceptibility to online fraud. Such computers can also become inadvertent participants in or components of an online crime network, spam network, and/or phishing network as well as be used as a part of a distributed denial-of-service attack. Mitigating the effects of and remediating the installations of malicious bots will make it more difficult for botnets to operate and could reduce the level of online crime on the Internet in general and/or on a particular Internet Service Provider's network.

1.1. Key Terminology

This section defines the key terms used in this document.

1.1.1. Malicious Bots, or Bots

A malicious or potentially malicious bot (derived from the word "robot", hereafter simply referred to as a "bot") refers to a program that is installed on a system in order to enable that system to automatically (or semi-automatically) perform a task or set of tasks typically under the command and control of a remote administrator, or "bot master". Bots are also known as "zombies". Such bots may have been installed surreptitiously, without the user's full understanding of what the bot will do once installed, unknowingly as part of another software installation, under false pretenses, and/or in a variety of other possible ways.

It is important to note that there are "good" bots. Such good bots are often found interacting with a computing resource in environments such as gaming and Internet Relay Chat (IRC) [RFC1459], where a continual, interactive presence can be a requirement for participating in the games. Since such good bots are performing useful, lawful, and non-disruptive functions, there is no reason for a provider to monitor for their presence and/or alert users to their presence.

While there may be good, or harmless bots, for the purposes of this document, all mention of bots shall assume that the bots involved are malicious or potentially malicious in nature. Such malicious bots shall generally be assumed to have been deployed without the permission or conscious understanding of a particular Internet user. Thus, without a user's knowledge, bots may transform the user's

computing device into a platform from which malicious activities can be conducted. In addition, included explicitly in this category are potentially malicious bots, which may initially appear neutral but may simply be waiting for remote instructions to transform and/or otherwise begin engaging in malicious behavior. In general, installation of a malicious bot without user knowledge and consent is considered in most regions to be unlawful, and the activities of malicious bots typically involve unlawful or other maliciously disruptive activities.

1.1.2. Bot Networks, or Botnets

A "bot network", or "botnet", is defined as a concerted network of bots capable of acting on instructions generated remotely. The malicious activities are either focused on the information on the local machine or acting to provide services for remote machines. Bots are highly customizable so they can be programmed to do many things. The major malicious activities include but are not limited to identity theft, spam, spim (spam over Instant Messaging (IM)), spit (spam over Internet telephony), email address harvesting, distributed denial-of-service (DDoS) attacks, key-logging, fraudulent DNS pharming (redirection), hosting proxy services, fast flux (see Section 1.1.5) hosting, hosting of illegal content, use in man-in-the-middle attacks, and click fraud.

Infection vectors (infection pathways) include un-patched operating systems, software vulnerabilities (which include so-called zero-day vulnerabilities where no patch yet exists), weak/non-existent passwords, malicious web sites, un-patched browsers, malware, vulnerable helper applications, inherently insecure protocols, protocols implemented without security features switched on, and social engineering techniques to gain access to the user's computer. The detection and destruction of bots is an ongoing issue and also a constant battle between the Internet security community and network security engineers on the one hand and bot developers on the other.

Initially, some bots used IRC to communicate but were easy to shut down if the command and control server was identified and deactivated. Newer command and control methods have evolved, such that those currently employed by bot masters make them much more resistant to deactivation. With the introduction of peer-to-peer (P2P) architectures and associated protocols, the use of HTTP and other resilient communication protocols, and the widespread adoption of encryption, bots are considerably more difficult to identify and isolate from typical network usage. As a result, increased reliance is being placed on anomaly detection and behavioral analysis, both locally and remotely, to identify bots.

1.1.3. Host

As used in the context of this document, the host or computer of an end user is intended to refer to a computing device that connects to the Internet. This encompasses devices used by Internet users such as personal computers (including laptops, desktops, and netbooks), mobile phones, smart phones, home gateway devices, and other end user computing devices that are connected or can connect to the public Internet and/or private IP networks.

Increasingly, other household systems and devices contain embedded hosts that are connected to or can connect to the public Internet and/or private IP networks. However, these devices may not be under interactive control of the Internet user, such as may be the case with various smart home and smart grid devices.

1.1.4. Malware

Malware is short for "malicious software". In this case, malicious bots are considered a subset of malware. Other forms of malware could include viruses and other similar types of software. Internet users can sometimes cause their hosts to be infected with malware, which may include a bot or cause a bot to install itself, via inadvertently accessing a specific web site, downloading a file, or other activities.

In other cases, Internet-connected hosts may become infected with malware through externally initiated malicious activities such as the exploitation of vulnerabilities or the brute force guessing of access credentials.

1.1.5. Fast Flux

Domain Name System (DNS) fast fluxing occurs when a domain is bound in DNS using A records to multiple IP addresses, each of which has a very short Time-to-Live (TTL) value associated with it. This means that the domain resolves to varying IP addresses over a short period of time.

DNS fast flux is typically used in conjunction with proxies that are normally run on compromised user hosts. These proxies route the web requests to the real host, which serves the data being sought. The effect of this is to make the detection of the real host much more difficult and to ensure that the backend or hidden site remains up for as long as possible.

2. Problem Statement

Hosts used by Internet users, which in this case are customers of an Internet Service Provider (ISP), can be infected with malware that may contain and/or install one or more bots on a host. They can present a major problem for an ISP for a number of reasons (not to mention, of course, the problems created for users). First, these bots can be used to send spam, in some cases very large volumes of spam [Spamalytics]. This spam can result in extra cost for the ISPs in terms of wasted network, server, and/or personnel resources, among many other potential costs and side effects. Such spam can also negatively affect the reputation of the ISP, their customers, and the email reputation of the IP address space used by the ISP (often referred to simply as "IP reputation"). A further potential complication is that IP space compromised by bad reputation may continue to carry this bad reputation even when used for entirely innocent purposes following reassignment of that IP space.

In addition, these bots can act as platforms for directing, participating in, or otherwise conducting attacks on critical Internet infrastructure [Threat-Report]. Bots are frequently used as part of coordinated DDoS attacks for criminal, political, or other motivations [Gh0st][Dragon][DDoS]. For example, bots have been used to attack Internet resources and infrastructure ranging from web sites to email servers and DNS servers, as well as the critical Internet infrastructure of entire countries [Estonia][Combat-Zone]. Motivations for such coordinated DDoS attacks can range from criminal extortion attempts through to online protesting and nationalistic fervor [Whiz-Kid]. DDoS attacks may also be motivated by simple personal vendettas or by persons simply seeking a cheap thrill at the expense of others.

There is good evidence to suggest that bots are being used in the corporate environment for purposes of corporate espionage including the exfiltration of corporate financial data and intellectual property. This also extends to the possibility of bots being used for state-sponsored purposes such as espionage.

While any computing device can be infected with bots, the majority of bot infections affect the personal computers used by Internet end users. As a result of the role of ISPs in providing IP connectivity, among many other services, to Internet users, these ISPs are in a unique position to be able to attempt to detect and observe botnets operating in their networks. Furthermore, ISPs may also be in a unique position to be able to notify their customers of actual, potential, or likely infection by bots or other infection.

From the perspective of end users, being notified that they may have an infected computer on their network is important information. Once they know this, they can take steps to remove the bots, resolve any problems that may stem from the bot infection, and protect themselves against future threats. It is important to notify users that they may be infected with a bot because bots can consume vast amounts of local computing and network resources, enable theft of personal information (including personal financial information), enable the host to be used for criminal activities (that may result in the Internet user being legally culpable), and destroy or leave the host in an unrecoverable state via "kill switch" bot technologies.

As a result, the intent of this document is to provide guidance to ISPs and other organizations for the remediation of hosts infected with bots, so as to reduce the size of botnets and minimize the potential harm that bots can inflict upon Internet infrastructure in general as well as on individual Internet users. Efforts by ISPs and other organizations can, over time, reduce the pool of hosts infected with bots on the Internet, which in turn could result in smaller botnets with less capability for disruption.

The potential mitigation of bots is accomplished through a process of detection, notification to Internet users, and remediation of bot infections with a variety of tools, as described later in this document.

3. Important Notice of Limitations and Scope

The techniques described in this document in no way guarantee the remediation of all bots. Bot removal is potentially a task requiring specialized knowledge, skills, and tools; it may be beyond the ability of average users. Attempts at bot removal may frequently be unsuccessful, or only partially successful, leaving the user's system in an unstable and unsatisfactory state or even in a state where it is still infected. Attempts at bot removal can result in side effects ranging from a loss of data to partial or complete loss of system usability.

In general, the only way a user can be sure they have removed some of today's increasingly sophisticated malware is by "nuking-and-paving" the system: reformatting the drive, reinstalling the operating system and applications (including all patches) from scratch, and then restoring user files from a known clean backup. However, the introduction of persistent memory-based malware may mean that, in some cases, this may not be enough and may prove to be more than any end user can be reasonably expected to resolve [BIOS]. Experienced users would have to re-flash or re-image persistent memory sections or components of their hosts in order to remove persistent memory-

based malware. However, in some cases, not even nuking-and-paving the system will solve the problem, which calls for hard drive replacement and/or complete replacement of the host.

Devices with embedded operating systems, such as video gaming consoles and smart home appliances, will most likely be beyond a user's capability to remediate by themselves and could therefore require the aid of vendor-specific advice, updates, and tools. However, in some cases, such devices will have a function or switch to enable the user to reset that device to a factory default configuration, which may sometimes enable the user to remediate the infection. Care should be taken when imparting remediation advice to Internet users given the increasingly wide array of computing devices that can be, or could be, infected by bots in the future.

This document is not intended to address the issues relating to the prevention of bots on an end user device. This is out of the scope of this document.

4. Detection of Bots

An ISP must first identify that an Internet user is infected or likely to have been infected with a bot (a user is assumed to be their customer or otherwise connected to the ISP's network). The ISP should attempt to detect the presence of bots using methods, processes, and tools that maintain the privacy of the personally identifiable information (PII) of their customers. The ISP should not block legitimate traffic in the course of bot detection and should instead employ detection methods, tools, and processes that seek to be non-disruptive and transparent to Internet users and end user applications.

Detection methods, tools, and processes may include analysis of specific network and/or application traffic flows (such as traffic to an email server), analysis of aggregate network and/or application traffic data, data feeds received from other ISPs and organizations (such as lists of the ISP's IP addresses that have been reported to have sent spam), feedback from the ISP's customers or other Internet users, as well as a wide variety of other possibilities. In practice, it has proven effective to confirm a bot infection through the use of a combination of multiple bot detection data points. This can help to corroborate information of varying dependability or consistency, as well as to avoid or minimize the possibility of false positive identification of hosts. Detection should also, where possible and feasible, attempt to classify the specific bot infection type in order to confirm that it is malicious in nature, estimate the variety and severity of threats it may pose (such as spam bot, key-logging bot, file distribution bot, etc.), and determine potential

methods for eventual remediation. However, given the dynamic nature of botnet management and the criminal incentives to seek quick financial rewards, botnets frequently update or change their core capabilities. As a consequence, botnets that are initially detected and classified by the ISP as made up of one particular type of bot need to be continuously monitored and tracked in order to correctly identify the threat the botnet poses at any particular point in time.

Detection is also time sensitive. If complex analysis is required and multiple confirmations are needed to verify a bot is indeed present, then it is possible that the bot may cause some damage (to either the infected host or a remotely targeted system) before it can be stopped. This means that an ISP needs to balance the desire or need to definitively classify and/or confirm the presence of a bot, which may take an extended period of time, with the ability to predict the likelihood of a bot in a very short period of time. Such determinations must have a relatively low false positive rate in order to maintain the trust of users. This "definitive-versus-likely" challenge is difficult and, when in doubt, ISPs should err on the side of caution by communicating that a bot infection has taken place. This also means that Internet users may benefit from the installation of client-based security software on their host. This can enable rapid heuristically based detection of bot activity, such as the detection of a bot as it starts to communicate with other botnets and execute commands. Any bot detection system should also be capable of adapting, either via manual intervention or automatically, in order to cope with a rapidly evolving threat.

As noted above, detection methods, tools, and processes should ensure that privacy of customers' personally identifiable information (PII) is maintained. This protection afforded to PII should also extend to third parties processing data on behalf of ISPs. While bot detection methods, tools, and processes are similar to spam and virus defenses deployed by the ISP for the benefit of their customers (and may be directly related to those defenses), attempts to detect bots should take into account the need of an ISP to take care to ensure any PII collected or incidentally detected is properly protected. This is important because just as spam defenses may involve scanning the content of email messages, which may contain PII, then so too may bot defenses similarly come into incidental contact with PII. The definition of PII varies from one jurisdiction to the next so proper care should be taken to ensure that any actions taken comply with legislation and good practice in the jurisdiction in which the PII is gathered. Finally, depending upon the geographic region within which an ISP operates, certain methods relating to bot detection may need to be included in relevant terms of service documents or other documents that are available to the customers of a particular ISP.

There are several bot detection methods, tools, and processes that an ISP may choose to utilize, as noted in the list below. It is important to note that the technical solutions available are relatively immature and are likely to change over time, evolving rapidly in the coming years. While these items are described in relation to ISPs, they may also be applicable to organizations operating other networks, such as campus networks and enterprise networks.

- a. Where it is not legally proscribed and an accepted industry practice in a particular market region, an ISP may in some manner "scan" its IP space in order to detect un-patched or otherwise vulnerable hosts or to detect the signs of infection. This may provide the ISP with the opportunity to easily identify Internet users who appear already to be infected or are at great risk of being infected with a bot. ISPs should note that some types of port scanning may leave network services in a hung state or render them unusable due to common frailties and that many modern firewall and host-based intrusion detection implementations may alert the Internet user to the scan. As a result, the scan may be interpreted as a malicious attack against the host. Vulnerability scanning has a higher probability of leaving accessible network services and applications in a damaged state and will often result in a higher probability of detection by the Internet user and subsequent interpretation as a targeted attack. Depending upon the vulnerability for which an ISP may be scanning, some automated methods of vulnerability checking may result in data being altered or created afresh on the Internet user's host, which can be a problem in many legal environments. It should also be noted that due to the prevalence of Network Address Translation devices, Port Address Translation devices, and/or firewall devices in user networks, network-based vulnerability scanning may be of limited value. Thus, while we note that this is one technique that may be utilized, it is unlikely to be particularly effective and has problematic side effects, which leads the authors to recommend against the use of this particular method.
- b. An ISP may also communicate and share selected data, via feedback loops or other mechanisms, with various third parties. Feedback loops are consistently formatted feeds of real-time (or nearly real-time) abuse reports offered by threat data clearinghouses, security alert organizations, other ISPs, and other organizations. The formats for feedback loops include those defined in both the Abuse Reporting Format (ARF) [RFC5965] and the Incident Object Description Exchange Format (IODEF) [RFC5070]. The data may include, but is not limited to, IP addresses of hosts that appear to be either definitely or

probably infected, IP addresses, domain names or fully qualified domain names (FQDNs) known to host malware and/or be involved in the command and control of botnets, recently tested or discovered techniques for detecting or remediating bot infections, new threat vectors, and other relevant information. A few good examples of data sharing are noted in Appendix A.

- c. An ISP may use Netflow [RFC3954] or other similar passive network monitoring to identify network anomalies that may be indicative of botnet attacks or bot communications. For example, an ISP may be able to identify compromised hosts by identifying traffic destined to IP addresses associated with the command and control of botnets or destined to the combination of an IP address and control port associated with a command and control network (sometimes command and control traffic comes from a host that has legitimate traffic). In addition, bots may be identified when a remote host is under a DDoS attack, because hosts participating in the attack will likely be infected by a bot. This can often be observed at network borders although ISPs should beware of source IP address spoofing techniques that may be employed to avoid or confuse detection.
- d. An ISP may use DNS-based techniques to perform detection. For example, a given classified bot may be known to query a specific list of domain names at specific times or on specific dates (in the example of the so-called "Conficker" bot (see [Conficker]), often by matching DNS queries to a well-known list of domains associated with malware. In many cases, such lists are distributed by or shared using third parties, such as threat data clearinghouses.
- e. Because hosts infected by bots are frequently used to send spam or participate in DDoS attacks, the ISP servicing those hosts will normally receive complaints about the malicious network traffic. Those complaints may be sent to role accounts specified in RFC 2142 [RFC2142], such as abuse@, or to other relevant addresses such as to abuse or security addresses specified by the site as part of its WHOIS (or other) contact data.
- f. ISPs may also discover likely bot-infected hosts located on other networks. Thus, when legally permissible in a particular market region, it may be worthwhile for ISPs to share information relating to those compromised hosts with the relevant remote network operator, security researchers, and blocklist operators.

- g. ISPs may operate or subscribe to services that provide "sinkholing" or "honeynet" capabilities. This may enable the ISP to obtain near-real-time lists of bot-infected hosts as they attempt to join a larger botnet or propagate to other hosts on a network.
- h. ISP industry associations should examine the possibility of collating statistics from ISP members in order to provide good statistics about bot infections based on real ISP data.
- i. An Intrusion Detection System (IDS) can be a useful tool to actually help identify the malware. An IDS tool such as Snort (open source IDS platform; see [Snort]) can be placed in a walled garden and used to analyze end user traffic to confirm malware type. This will help with remediation of the infected device.

5. Notification to Internet Users

Once an ISP has detected a bot, or the strong likelihood of a bot, steps should be undertaken to inform the Internet user that they may have a bot-related problem. An ISP should decide the most appropriate method or methods for providing notification to one or more of their customers or Internet users, depending upon a range of factors including the technical capabilities of the ISP, the technical attributes of its network, financial considerations, available server resources, available organizational resources, the number of likely infected hosts detected at any given time, and the severity of any possible threats. Such notification methods may include one or more of the methods described in the following subsections, as well as other possible methods not described below.

It is important to note that none of these methods are guaranteed to be one hundred percent successful and that each has its own set of limitations. In addition, in some cases, an ISP may determine that a combination of two or more methods is most appropriate and effective and reduces the chance that malware may block a notification. As such, the authors recommend the use of multiple notification methods. Finally, notification is also considered time sensitive; if the user does not receive or view the notification in a timely fashion, then a particular bot could launch an attack, exploit the user, or cause other harm. If possible, an ISP should establish a preferred means of communication when the subscriber first signs up for service. As a part of the notification process, ISPs should maintain a record of the allocation of IP addresses to subscribers for a period long enough to allow any commonly used bot detection technology to be able to accurately link an infected IP address to a subscriber. This

record should only be maintained for a period of time that is necessary to support bot detection, but no longer, in order to protect the privacy of the individual subscriber.

One important factor to bear in mind is that notification to end users needs to be resistant to potential spoofing. This should be done to protect, as reasonably as possible, against the potential of legitimate notifications being spoofed and/or used by parties with intent to perform additional malicious attacks against victims of malware or even to deliver additional malware.

It should be possible for the end user to indicate the preferred means of notification on an opt-in basis for that notification method. It is recommended that the end user should not be allowed to opt out of notification entirely.

When users are notified, an ISP should endeavor to give as much information as possible to the end user regarding which bot detection methods are employed at the ISP, consonant with not providing information to those creating or deploying the bots so that they would be able to avoid detection.

5.1. Email Notification

This is a common form of notification used by ISPs. One drawback of using email is that it is not guaranteed to be viewed within a reasonable time frame, if at all. The user may be using a different primary email address than the one they provided to the ISP. In addition, some ISPs do not provide an email account at all as part of a bundle of Internet services and/or do not have a need for or method by which to request or retain the primary email addresses of Internet users of their networks. Another possibility is that the user, their email client, and/or their email servers could determine or classify such a notification as spam, which could delete the message or otherwise file it in an email folder that the user may not check on a regular and/or timely basis. Bot masters have also been known to impersonate the ISP or trusted sender and send fraudulent emails to the users. This technique of social engineering often leads to new bot infestations. Finally, if the user's email credentials are compromised, then a hacker and/or a bot could simply access the user's email account and delete the email before it is read by the user.

5.2. Telephone Call Notification

A telephone call may be an effective means of communication in particularly high-risk situations. However, telephone calls may not be feasible due to the cost of making a large number of calls, as

measured in either time, money, organizational resources, server resources, or some other means. In addition, there is no guarantee that the user will answer their phone. To the extent that the telephone number called by the ISP can be answered by the infected computing device, the bot on that host may be able to disconnect, divert, or otherwise interfere with an incoming call. Users may also interpret such a telephone notification as a telemarketing call and therefore not welcome it or not accept the call at all. Finally, even if a representative of the ISP is able to connect with and speak to a user, that user is very likely to lack the necessary technical expertise to understand or be able to effectively deal with the threat.

5.3. Postal Mail Notification

This form of notification is probably the least popular and effective means of communication, due to preparation time, delivery time, the cost of printing and paper, and the cost of postage.

5.4. Walled Garden Notification

Placing a user in a walled garden is another approach that ISPs may take to notify users. A "walled garden" refers to an environment that controls the information and services that a subscriber is allowed to utilize and what network access permissions are granted. A walled garden implementation can range from strict to leaky. In a strict walled garden environment, access to most Internet resources is typically limited by the ISP. In contrast, a leaky walled garden environment permits access to all Internet resources, except those deemed malicious, and ensures access to those that can be used to notify users of infections.

Walled gardens are effective because it is possible to notify the user and simultaneously block all communication between the bot and the command and control channel. While in many cases the user is almost guaranteed to view the notification message and take any appropriate remediation actions, this approach can pose other challenges. For example, it is not always the case that a user is actively utilizing a host that implements a web browser, has a web browser actively running on it, or operates another application that uses ports that are redirected to the walled garden. In one example, a user could be playing a game online, via the use of a dedicated, Internet-connected game console. In another example, the user may not be using a host with a web browser when they are placed in the walled garden and may instead be in the course of a telephone conversation or may be expecting to receive a call using a Voice over IP (VoIP) device of some type. As a result, the ISP may feel the need to maintain a potentially lengthy white list of domains that are

not subject to the typical restrictions of a walled garden, which could well prove to be an onerous task from an operational perspective.

For these reasons, the implementation of a leaky walled garden makes more sense, but a leaky walled garden has a different set of drawbacks. The ISP has to assume that the user will eventually use a web browser to acknowledge the notification; otherwise, the user will remain in the walled garden and not know it. If the intent of the leaky walled garden is solely to notify the user about the bot infection, then the leaky walled garden is not ideal because notification is time sensitive, and the user may not receive the notification until the user invokes a request for the targeted service and/or resource. This means the bot can potentially do more damage. Additionally, the ISP has to identify which services and/or resources to restrict for the purposes of notification. This does not have to be resource specific and can be time based and/or policy based. An example of how notification could be made on a timed basis could involve notification for all HTTP requests every 10 minutes, or show the notification for one in five HTTP requests.

The ISP has several options to determine when to let the user out of the walled garden. One approach may be to let the user determine when to exit. This option is suggested when the primary purpose of the walled garden is to notify users and provide information on remediation only, particularly since notification is not a guarantee of successful remediation. It could also be the case that, for whatever reason, the user makes the judgment that they cannot then take the time to remediate their host and that other online activities that they would like to resume are more important. Exit from the walled garden may also involve a process to verify that it is indeed the user who is requesting exit from the walled garden and not the bot.

Once the user acknowledges the notification, they may decide either to remediate and exit the walled garden or to exit the walled garden without remediating the issue. Another approach may be to enforce a stricter policy and require the user to clean the host prior to permitting the user to exit the walled garden, though this may not be technically feasible depending upon the type of bot, obfuscation techniques employed by a bot, and/or a range of other factors. Thus, the ISP may also need to support tools to scan the infected host (in the style of a virus scan, rather than a port scan) and determine whether it is still infected or rely on user judgment that the bot has been disabled or removed. One challenge with this approach is that the user might have multiple hosts sharing a single IP address, such as via a common home gateway device that performs Network

Address Translation (NAT). In such a case, the ISP may need to determine from user feedback, or other means, that all affected hosts have been remediated, which may or may not be technically feasible.

Finally, when a walled garden is used, a list of well-known addresses for both operating system vendors and security vendors should be created and maintained in a white list that permits access to these sites. This can be important for allowing access from the walled garden by end users in search of operating system and application patches. It is recommended that walled gardens be seriously considered as a method of notification as they are easy to implement and proven to be effective as a means of getting end user attention.

5.5. Instant Message Notification

IM provides the ISP with a simple means to communicate with the user. There are several advantages to using IM that make it an attractive option. If the ISP provides IM service and the user subscribes to it, then the user can be notified easily. IM-based notification can be a cost-effective means to communicate with users automatically from an IM alert system or by a manual process, involving the ISP's support staff. Ideally, the ISP should allow the user to register their IM identity in an ISP account management system and grant permission to be contacted via this means. If the IM service provider supports off-line messaging, then the user can be notified regardless of whether they are currently logged into the IM system.

There are several drawbacks with this communications method. There is a high probability that a subscriber may interpret the communication to be spam and thus ignore it. Also, not every user uses IM and/or the user may not provide their IM identity to the ISP so some alternative means have to be used. Even in those cases where a user does have an IM address, they may not be signed onto that IM system when the notification is attempted. There may be a privacy concern on the part of users when such an IM notification must be transmitted over a third-party network and/or IM service. As such, should this method be used, the notification should be discreet and not include any PII in the notification itself.

5.6. Short Message Service (SMS) Notification

SMS allows the ISP to send a brief description of the problem to notify the user of the issue, typically to a mobile device such as a mobile phone or smart phone. Ideally, the ISP should allow the user to register their mobile number and/or SMS address in an ISP account management system and grant permission to be contacted via this means. The primary advantage of SMS is that users are familiar with

receiving text messages and are likely to read them. However, users may not act on the notification immediately if they are not in front of their host at the time of the SMS notification.

One disadvantage is that ISPs may have to follow up with an alternate means of notification if not all of the necessary information may be conveyed in one message, given constraints on the number of characters in an individual message (typically 140 characters). Another disadvantage with SMS is the cost associated with it. The ISP has to either build its own SMS gateway to interface with the various wireless network service providers or use a third-party SMS clearinghouse (relay) to notify users. In both cases, an ISP may incur fees related to SMS notifications, depending upon the method used to send the notifications. An additional downside is that SMS messages sent to a user may result in a charge to the user by their wireless provider, depending upon the plan to which they subscribe and the country in which the user resides. Another minor disadvantage is that it is possible to notify the wrong user if the intended user changes their mobile number but forgets to update it with the ISP.

There are several other drawbacks with this communications method. There is a high probability that subscriber may interpret the communication to be spam and thus ignore it. Also, not every user uses SMS, and/or the user may not provide their SMS address or mobile number to the ISP. Even in those cases where a user does have an SMS address or mobile number, their device may not be powered on or otherwise available on a wireless network when the notification is attempted. There may also be a privacy concern on the part of users when such an SMS notification must be transmitted over a third-party network and/or SMS clearinghouse. As such, should this method be used, the notification should be discreet and not include any PII in the notification itself.

5.7. Web Browser Notification

Near-real-time notification to the user's web browser is another technique that may be utilized for notifying the user [RFC6108], though how such a system might operate is outside the scope of this document. Such a notification could have a comparative advantage over a walled garden notification, in that it does not restrict traffic to a specified list of destinations in the same way that a walled garden would, by definition. However, as with a walled garden notification, there is no guarantee that a user is making use of a web browser at any given time, though such a system could certainly provide a notification when such a browser is eventually used. Compared to a walled garden, a web browser notification is probably

preferred from the perspective of Internet users, as it does not have the risk of disrupting non-web sessions, such as online games, VoIP calls, etc. (as noted in Section 5.4).

There are alternative methods of web browser notification offered commercially by a number of vendors. Many of the techniques used are proprietary, and it is not within the scope of this document to describe how they are implemented. These techniques have been successfully implemented at several ISPs.

It should be noted that web notification is only intended to notify devices running a web browser.

5.8. Considerations for Notification to Public Network Locations

Delivering a notification to a location that provides a shared public network, such as a train station, public square, coffee shop, or similar location may be of low value since the users connecting to such networks are typically highly transient and generally not known to site or network administrators. For example, a system may detect that a host on such a network has a bot, but by the time a notification is generated, that user has departed from the network and moved elsewhere.

5.9. Considerations for Notification to Network Locations Using a Shared IP Address

Delivering a notification to a location that accesses the Internet routed through one or more shared public IP addresses may be of low value since it may be quite difficult to differentiate between users when providing a notification. For example, on a business network of 500 users, all sharing one public IP address, it may be sub-optimal to provide a notification to all 500 users if you only need one specific user to be notified and take action. As a result, such networks may find value in establishing a localized bot detection and notification system, just as they are likely to also establish other localized systems for security, file sharing, email, and so on.

However, should an ISP implement some form of notification to such networks, it may be better to simply send notifications to a designated network administrator at the site. In such a case, the local network administrator may like to receive additional information in such a notification, such as a date and timestamp, the source port of the infected system, and malicious sites and ports that may have been visited.

5.10. Notification and End User Expertise

The ultimate effectiveness of any of the aforementioned forms of notification is heavily dependent upon both the expertise of the end user and the wording of any such notification. For example, while a user may receive and acknowledge a notification, that user may lack the necessary technical expertise to understand or be able to deal effectively with the threat. As a result, it is important that such notifications use clear and easily understood language, so that the majority of users (who are non-technical) may understand the notification. In addition, a notification should provide easily understood guidance on how to remediate a threat as described in Section 6, potentially with one path for technical users to take and another for non-technical users.

6. Remediation of Hosts Infected with a Bot

This section covers the different options available to remediate a host, which means to remove, disable, or otherwise render a bot harmless. Prior to this step, an ISP has detected the bot, notified the user that one of their hosts is infected with a bot, and now may provide some recommended means to clean the host. The generally recommended approach is to provide the necessary tools and education to the user so that they may perform bot remediation themselves, particularly given the risks and difficulties inherent in attempting to remove a bot.

For example, this may include the creation of a special web site with security-oriented content that is dedicated for this purpose. This should be a well-publicized security web site to which a user with a bot infection can be directed to for remediation. This security web site should clearly explain why the user was notified and may include an explanation of what bots are and the threats that they pose. There should be a clear explanation of the steps that the user should take in order to attempt to clean their host and information on how users can keep the host free of future infections. The security web site should also have a guided process that takes non-technical users through the remediation process, on an easily understood, step-by-step basis.

In terms of the text used to explain what bots are and the threats that they pose, something simple such as this may suffice:

What is a bot? A bot is a piece of software, generally installed on your machine without your knowledge, which either sends spam or tries to steal your personal information. They can be very difficult to spot, though you may have noticed that your computer is running much more slowly than usual or you may notice regular

disk activity even when you are not doing anything. Ignoring this problem is risky to you and your personal information. Thus, bots need to be removed to protect your data and your personal information.

Many bots are designed to work in a very stealthy manner, and as such, there may be a need to make sure that the Internet user understands the magnitude of the threat faced despite the stealthy nature of the bot.

It is also important to note that it may not be immediately apparent to the Internet user precisely which devices have been infected with a particular bot. This may be due to the user's home network configuration, which may encompass several hosts, where a home gateway that performs Network Address Translation (NAT) to share a single public IP address has been used. Therefore, any of these devices can be infected with a bot. The consequence of this for an ISP is that remediation advice may not ultimately be immediately actionable by the Internet user, as that user may need to perform additional investigation within their own home network.

An added complication is that the user may have a bot infection on a device such as a video console, multimedia system, appliance, or other end user computing device that does not have a typical desktop computing interface. As a result, diligence needs to be taken by the ISP where possible such that it can identify and communicate the specific nature of the device that has been infected with a bot and provide further appropriate remediation advice. If the ISP cannot pin down the device or identify its type, then it should make it clear to the user that any initial advice given is generic and further advice can be given (or is available) once the type of infected device is known.

There are a number of forums that exist online to provide security-related support to end users. These forums are staffed by volunteers and often are focused around the use of a common tool set to help end users to remediate hosts infected with malware. It may be advantageous to ISPs to foster a relationship with one or more forums, perhaps by offering free hosting or other forms of sponsorship.

It is also important to keep in mind that not all users will be technically adept, as noted in Section 5.10. As a result, it may be more effective to provide a range of suggestion options for remediation. This may include, for example, a very detailed "do it yourself" approach for experts, a simpler guided process for the average user, and even assisted remediation as described in Section 6.2.

6.1. Guided Remediation Process

Minimally, the Guided Remediation Process should include the following goals, with options and/or recommendations for achieving them:

1. Back up personal files. For example:

Before you start, make sure to back up all of your important data. (You should do this on a regular basis anyway.) You can back up your files manually or using a system backup software utility, which may be part of your Operating System (OS). You can back up your files to a USB Thumb Drive (aka USB Key), a writable CD/DVD-ROM, an external hard drive, a network file server, or an Internet-based backup service.

It may be advisable to suggest that the user backup is performed onto separate backup media or devices if they suspect bot infection.

2. Download OS patches and Anti-Virus (A/V) software updates. For example, links could be provided to Microsoft Windows updates, Apple Mac OS updates, or other major operating systems that are relevant to users and their devices.
3. Configure the host to automatically install updates for the OS, A/V, and other common web browsers such as Microsoft Internet Explorer, Mozilla Firefox, Apple Safari, Opera, and Google Chrome.
4. Get professional assistance if they are unable to remove the bots themselves. If purchasing professional assistance, then the user should be encouraged to predetermine how much they are willing to pay for that help. For example, if the host that is being remediated is old and can easily be replaced with a new, faster, larger, and more reliable system for a certain cost, then it makes no sense to spend more than that cost to fix the old host. On the other hand, if the customer has a brand-new host, it might make perfect sense to spend the money to attempt to remediate it.
5. To continue, regardless of whether the user or a knowledgeable technical assistant is working on remediating the host, the first task should be to determine which of multiple potentially infected machines may be the one that needs attention (in the common case of multiple hosts in a home network). Sometimes, as in cases where there is only a single directly attached host, or the user has been noticing problems with one of their hosts, this can be easy. Other times, it may be more difficult, especially

if there are no clues as to which host is infected. If the user is behind a home gateway/router, then the first task may be to ascertain which of the machines is infected. In some cases, the user may have to check all machines to identify the infected one.

6. ISPs may also look at offering a CD/DVD with remediation processes and software in the event that a host is so badly infected as to be unable to communicate over the Internet.
 7. User surveys to solicit feedback on whether the notification and remediation process is effective and what recommended changes could be made in order to improve the ease, understandability, and effectiveness the remediation process.
 8. If the user is interested in reporting the host's bot infection to an applicable law enforcement authority, then the host effectively becomes a cyber "crime scene", and the infection should not be mitigated unless or until law enforcement has collected the necessary evidence. For individuals in this situation, the ISP may wish to provide links to local, state, federal, or other relevant computer crime offices. (Note: Some "minor" incidents, even if highly traumatic to the user, may not be sufficiently serious for law enforcement to commit some of their limited resources to an investigation.) In addition, individual regions may have other, specialized computer crime organizations to which these incidents can be reported. For example, in the United States, that organization is the Internet Crime Complaint Center, at <http://www.ic3.gov>.
 9. Users may also be interested in links to security expert forums, where other users can assist them.
- 6.2. Professionally Assisted Remediation Process

It should be acknowledged that, based on the current state of remediation tools and the technical abilities of end users, that many users may be unable to remediate on their own. As a result, it is recommended that users have the option for professional assistance. This may entail online or telephone assistance for remediation, as well as working face to face with a professional who has training and expertise in the removal of malware. It should be made clear at the time of offering this service that this service is intended for those that do not have the skills or confidence to attempt remediation and is not intended as an up-sell by the ISP.

7. Failure or Refusal to Remediate

ISP systems should track the bot infection history of hosts in order to detect when users consistently fail to remediate or refuse to take any steps to remediate. In such cases, ISPs may need to consider taking additional steps to protect their network, other users and hosts on that network, and other networks. Such steps may include a progression of actions up to and including account termination. Refusal to remediate can be viewed as a business issue, and as such, no technical recommendation is possible.

8. Sharing of Data from the User to the ISP

As an additional consideration, it may be useful to create a process by which users could choose, at their option and with their express consent, to share data regarding their bot infections with their ISP and/or another authorized third party. Such third parties may include governmental entities that aggregate threat data, such as the Internet Crime Complaint Center referred to earlier in this document, academic institutions, and/or security researchers. While in many cases the information shared with the user's ISP or designated third parties will only be used for aggregated statistical analysis, it is also possible that certain research needs may be best met with more detailed data. Thus, any such data sharing from a user to the ISP or authorized third party may contain some type of personally identifiable information, either by design or inadvertently. As a result, any such data sharing should be enabled on an opt-in basis, where users review and approve of the data being shared and the parties with which it is to be shared, unless the ISP is already required to share such data in order to comply with local laws and applicable regulations.

9. Security Considerations

This document describes in detail the numerous security risks and concerns relating to botnets. As such, it has been appropriate to include specific information about security in each section above. This document describes the security risks related to malicious bot infections themselves, such as enabling identity theft, theft of authentication credentials, and the use of a host to unwittingly participate in a DDoS attack, among many other risks. Finally, the document also describes security risks that may relate to the particular methods of communicating a notification to Internet users. Bot networks and bot infections pose extremely serious security risks, so readers should review this document carefully.

In addition, regarding notifications as described in Section 5, care should be taken to assure users that notifications have been provided by a trustworthy site and/or party, so that the notification is more difficult for phishers and/or malicious parties using social engineering tactics to mimic. Otherwise, care should be taken to ensure that the user has some level of trust that the notification is valid and/or that the user has some way to verify via some other mechanism or step that the notification is valid.

10. Privacy Considerations

This document describes at a high level the activities to which ISPs should be sensitive, i.e., where the collection or communication of PII may be possible. In addition, when performing notifications to end users (see Section 5), those notifications should not include PII.

As noted in Section 8, any sharing of data from the user to the ISP and/or authorized third parties should be done on an opt-in basis. Additionally the ISP and or authorized third parties should clearly state what data will be shared and with whom the data will be shared.

Lastly, as noted in other sections, there may be legal requirements in particular legal jurisdictions concerning how long any subscriber-related or other data is retained. An ISP operating in such a jurisdiction should be aware of these requirements and should comply with them.

11. Acknowledgements

The authors wish to acknowledge the following individuals and groups for performing a detailed review of this document and/or providing comments and feedback that helped to improve and evolve this document:

Mark Baugher

Richard Bennett

James Butler

Vint Cerf

Alissa Cooper

Jonathan Curtis

Jeff Chan

Roland Dobbins
Dave Farber
Stephen Farrell
Eliot Gillum
Joel Halpern
Joel Jaeggli
Scott Keoseyan
Murray S. Kucherawy
The Messaging Anti-Abuse Working Group (MAAWG)
Jose Nazario
Gunter Ollmann
David Reed
Roger Safian
Donald Smith
Joe Stewart
Forrest Swick
Sean Turner
Robb Topolski
Maxim Weinstein
Eric Ziegast

12. Informative References

- [BIOS] Sacco, A. and A. Ortega, "Persistent BIOS Infection", March 2009, <http://www.coresecurity.com/files/attachments/Persistent_BIOS_Infection_CanSecWest09.pdf>.
- [Combat-Zone] Alshech, E., "Cyberspace as a Combat Zone: The Phenomenon of Electronic Jihad", February 2007, <<http://www.memrijttm.org/content/en/report.htm?report=1822>>.
- [Conficker] Porras, P., Saidi, H., and V. Yegneswaran, "An Analysis of Conficker's Logic and Rendezvous Points", March 2009, <<http://mtc.sri.com/Conficker/>>.
- [DDoS] Saafan, A., "Distributed Denial of Service Attacks: Explanation, Classification and Suggested Solutions", March 2009, <www.exploit-db.com/download_pdf/14738/>.
- [Dragon] Nagaraja, S. and R. Anderson, "The snooping dragon: social-malware surveillance of the Tibetan movement", March 2009, <<http://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-746.pdf>>.
- [Estonia] Evron, G., "Battling Botnets and Online Mobs: Estonia's Defense Efforts during the Internet War", 2008, <<http://journal.georgetown.edu/wp-content/uploads/9.1-Evron.pdf>>.
- [Gh0st] Vallentin, M., Whiteaker, J., and Y. Ben-David, "The Gh0st in the Shell: Network Security in the Himalayas", February 2010, <<http://www.infowar-monitor.net/wp-content/uploads/2010/02/cs294-28-paper.pdf>>.
- [RFC1459] Oikarinen, J. and D. Reed, "Internet Relay Chat Protocol", RFC 1459, May 1993.
- [RFC2142] Crocker, D., "MAILBOX NAMES FOR COMMON SERVICES, ROLES AND FUNCTIONS", RFC 2142, May 1997.
- [RFC3954] Claise, B., "Cisco Systems NetFlow Services Export Version 9", RFC 3954, October 2004.
- [RFC5070] Danyliw, R., Meijer, J., and Y. Demchenko, "The Incident Object Description Exchange Format", RFC 5070, December 2007.

- [RFC5965] Shafranovich, Y., Levine, J., and M. Kucherawy, "An Extensible Format for Email Feedback Reports", RFC 5965, August 2010.
- [RFC6108] Chung, C., Kasyanov, A., Livingood, J., Mody, N., and B. Van Lieu, "Comcast's Web Notification System Design", RFC 6108, February 2011.
- [Snort] Roesch, M., "Snort Home Page", March 2009, <<http://www.snort.org/>>.
- [Spamalytics]
Kanich, C., Kreibich, C., Levchenko, K., Enright, B., Voelker, G., Paxson, V., and S. Savage, "Spamalytics: An Empirical Analysis of Spam Marketing Conversion", October 2008, <<http://www.icir.org/christian/publications/2008-ccs-spamalytics.pdf>>.
- [Threat-Report]
Ahamad, M., Amster, D., Barret, M., Cross, T., Heron, G., Jackson, D., King, J., Lee, W., Naraine, R., Ollman, G., Ramsey, J., Schmidt, H., and P. Traynor, "Emerging Cyber Threats Report for 2009: Data, Mobility and Questions of Responsibility will Drive Cyber Threats in 2009 and Beyond", October 2008, <<http://smartech.gatech.edu/bitstream/1853/26301/1/CyberThreatsReport2009.pdf>>.
- [Whiz-Kid] Berinato, S., "Case Study: How a Bookmaker and a Whiz Kid Took On a DDOS-based Online Extortion Attack", May 2005, <http://www.csoonline.com/article/220336/How_a_Bookmaker_and_a_Whiz_Kid_Took_On_a_DDOS_based_Online_Extortion_Attack>.

Appendix A. Examples of Third-Party Malware Lists

As noted in Section 4, there are many potential third parties that may be willing to share lists of infected hosts. This list is for example purposes only, is not intended to be either exclusive or exhaustive, and is subject to change over time.

- o Arbor - Atlas, see <http://atlas.arbor.net/>
- o Internet Systems Consortium - Secure Information Exchange (SIE), see <https://sie.isc.org/>
- o Microsoft - Smart Network Data Services (SNDS), see <https://postmaster.live.com/snds/>
- o SANS Institute / Internet Storm Center - DShield Distributed Intrusion Detection System, see <http://www.dshield.org/about.html>
- o ShadowServer Foundation, see <http://www.shadowserver.org/>
- o Spamhaus - Policy Block List (PBL), see <http://www.spamhaus.org/pbl/>
- o Spamhaus - Exploits Block List (XBL), see <http://www.spamhaus.org/xbl/>
- o Team Cymru - Community Services, see <http://www.team-cymru.org/>

Authors' Addresses

Jason Livingood
Comcast Cable Communications
One Comcast Center
1701 John F. Kennedy Boulevard
Philadelphia, PA 19103
USA

EEmail: jason_livingood@cable.comcast.com
URI: <http://www.comcast.com>

Nirmal Mody
Comcast Cable Communications
One Comcast Center
1701 John F. Kennedy Boulevard
Philadelphia, PA 19103
USA

EEmail: nirmal_mody@cable.comcast.com
URI: <http://www.comcast.com>

Mike O'Reirdan
Comcast Cable Communications
One Comcast Center
1701 John F. Kennedy Boulevard
Philadelphia, PA 19103
USA

EEmail: michael_oreirdan@cable.comcast.com
URI: <http://www.comcast.com>