To: National Institute of Standards and Technology, United States Department of Commerce
From: Messaging Anti-Abuse Working Group (MAAWG)
Date: November 14, 2011
Subject: MAAWG Comments on Models to Advance Voluntary Corporate Notification to Consumers Regarding the Illicit Use of Computer Equipment by Botnets and Related Malware

To whom it may concern:

Thank you for the opportunity to comment on the U.S. Department of Commerce (DoC) and U.S. Department of Homeland Security (DHS) "Models to Advance Voluntary Corporate Notification to Consumers Regarding the Illicit Use of Computer Equipment by Botnets and Related Malware," hereafter the RFI (Request for Information), (http://www.federalregister.gov/articles/2011/09/21/2011-24180/models-to-advance-voluntary-corporate-notification-to-consumers-regarding-the-illicit-use-of#p-6).

I. Introduction

The Messaging Anti-Abuse Working Group (MAAWG) is an international nonprofit, industry-led organization founded to fight online abuse such as phishing, botnets, fraud, spam, viruses and denial-of service attacks that can cause great harm to both individuals and national economies. MAAWG draws technical experts, researchers and policy specialists from a broad base of Internet service providers and network operators representing over one billion mailboxes, and from key technology providers, academia and volume sender organizations. The multi-disciplinary approach at MAAWG (http://www.maawg.org/) includes education, advice on public policy and legislation, development of industry best practices, guidance in the development of industry standards and the facilitation of global collaboration.

Our comments in this document generally follow the order in which topics were introduced in the RFI. In some cases, unless we have misunderstood the RFI author's intent, essentially the same question appears to have been asked multiple times. When we noticed this situation, we generally only responded once, referring the reader to our earlier response on the same or an extremely similar question.

We would also like to take this opportunity to remind the reader of the earlier submission by MAAWG, dated July 29th, 2011, to the Department of Commerce Internet Policy Task Force's green paper on "Cybersecurity, Innovation and the Internet Economy," www.maawg.org/sites/maawg/files/news/MAAWG_DoC_Internet_Task_Force-2011-08.pdf. We would like to incorporate and reiterate the points that we expressed in that document here. If there are perceptible or actual conflicts between the perspective expressed in this document and the earlier document, the positions expressed in the preceding document should be considered definitive, unless an expressed note disclaiming that earlier perspective is made herein.

II. The "Background" Section of the RFI

1. Overall Goal of the Department of Commerce/Department of Homeland Security Initiative

The RFI states:

> *Through this Request for Information and any follow-on work, the two Departments aim to reduce the harm that botnets inflict on the nation's computing environment.*

MAAWG strongly supports this overall goal and commends the Departments for undertaking this important inquiry. Combating the menace presented by bots is a top priority for MAAWG and we believe it should be a top cybersecurity priority for the DoC and DHS as well.

We may not always agree on all the precise details, but we definitely agree that working to reduce the prevalence and impact of botnets is critical to the security of the global Internet and all its users.

2. CSRIC's Best Practices for Botnet Protection and IETF Recommendations for Botnet Remediation

In providing background information and context for the inquiry, the RFI mentions two bodies of prior work which MAAWG agrees are particularly noteworthy, both of which MAAWG also strongly supports and endorses:

- The work of the FCC's Communications Security, Reliability and Interoperability Council (CSRIC) Working Group 8, which produced two dozen best practices addressing botnet protection for end users and networks (see the RFI at footnote 9), and

- The IETF draft "Recommendations for the Remediation of Bots in ISP Networks," most recently http://tools.ietf.org/html/draft-oreirdan-mody-bot-remediation-18 (expires April 28, 2012)

We would like to explicitly note and commend the hard work of all participants involved in this work, including the work and leadership of MAAWG-affiliated individuals. Notwithstanding these excellent prior efforts, obviously much work still remains to be done.

III. The "Incentives and Voluntary Approaches" Section of the RFI

3. Potential "Safe Harbor" Provisions

The RFI begins this section with the comment:

> *To promote voluntary best practices in botnet detection, notification and mitigation, one suggestion has been to provide companies that take action with certain types of liability protection in order to foster greater marketplace certainty.*

Liability worries, to the extent they exist, are a reflection of the litigious society in which American ISPs exist. An ISP, acting as a "good Samaritan" to voluntarily identify, notify and help a customer remediate infected systems may all-too-often find itself the target of customer ire, rather than customer appreciation:

- Users, not understanding the sometimes-proprietary technical approaches used to identify botnet-infected hosts, may view unexpected notifications about botnet infections as a potential invasive of their privacy or as intrusive. ("How do you know my computer is infected? Have you been spying on me? I'm going to sue you!")
- Notifications, unless persistent or done in a way that is impossible to ignore, may be disregarded; if done in an impossible-to-ignore way, users may resent those contacts. ("Quit bugging me or I'm going to call my attorney!")
- Efforts to help users attempt to remove malware from their system may subject an ISP intervener to putative liability for any subsequent issues that system may experience, whether related to the malware infection, the ISP's efforts to help remove that malware, or some wholly unrelated cause. ("My system has been screwed up ever since you tried to remove the malware you claimed was on it, and there's still something wrong with it! Thanks for nothing – I'll see you in court!")

Providing ISPs with limited immunity from legal liability for their voluntary efforts to combat bots has the potential to at least remove one potential obstacle to ISP efforts in that area – litigation exposure. As such, it would be welcomed, *provided* it does not inadvertently or unintentionally serve to provide disincentives for ISP intervention.

An example of liability protection that would be counterproductive and a disincentive to ISP action against bots might be offering ISPs the option of obtaining complete "common-carrier"-like immunity from liability whereby they only provide connectivity and have no responsibility (or even ability) to do anything beyond simply providing transport. Since this would discourage ISPs from proactively engaging in the fight against malware, MAAWG would oppose any such effort as ultimately ineffective and contrary to the best interests of its members and to the Internet as a whole. ISPs *must* remain involved and engaged in the war against bots.

Ultimately, however, even if a statute does protect an ISP from litigation, no statute can immunize an ISP from the potential economic impact of disgruntled customers "voting with their pocketbook." While a customer might be prevented from suing an ISP, most consumers have the option of moving to another provider if they are sufficiently dissatisfied with how their current ISP operates and interfaces with them, whether that is in conjunction with fighting an infection or an unrelated issue. Thus, even if an ISP is offered limited immunity from liability by statute, many ISPs may nonetheless be inclined toward employing a "light touch" in undertaking customer botnet identification, notification and remediation – unless all market participants are highly incented to meet a common standard.

Immunization from liability also does not address a related basic but critical issue: Who will pay? While limiting liability removes one potential financial obstacle (the risk of litigation), ISPs still face the actual operational costs associated with providing identification, notification and remediation services. On balance, the value of the proposed limits on liability may be small for most ISPs relative to the concrete and ongoing operational costs.

Thus, immunization from liability might best be viewed as helpful, or enabling. However, liability-related concerns are not the only obstacle that might deter some ISPs from proceeding with botnet identification, notification, and remediation tasks.

4. The "Centralized Consumer Resource Center" Concept

The RFI describes this notion as follows:

> *Another suggestion is to encourage ISPs to send consumer support queries to a centralized consumer resource center that could be supported by a wide number of players. Such a resource center could reduce the burden on corporate customer support centers by pooling resources. The center could aid consumers by, for example, providing certain no-cost means of support, as well as information on other means for expedited support. This center could also be used to facilitate information sharing and research that could lead to better botnet detection. Moreover, as a "condition of sponsorship" private sector entities could be required to adopt an agreed upon set of practices.*

The RFI then goes on to describe several potential models for that center (wholly private, mixed private/public, wholly public). In thinking about this concept, MAAWG has concerns about the extent to which this concept will be practical or scalable. Please remember:

- Many online technical self-help resources for botted-system owners already exist. Why would creating yet another such resource succeed if existing resources have not?
- Most botted systems are owned by average men and women, not technical enthusiasts. Extensive "hand holding" will likely be needed to help these users, not just a few quick tips. To get some idea of the magnitude of what might be required, consider the following scenario:
  - Assume each botted user requires an hour of a technician's time–a very conservative estimate that can easily increase by a factor of 2X or more higher.
  - While the RFI quotes a figure of 4,000,000 new botted hosts per month, let us assume there are only an average of 100,000 consumers of new U.S.-origin botted hosts per week that contact the proposed center.
  - To simplify for the sake of this model, assume that all technicians can be efficiently scheduled. (In reality, users would likely seek help from the center in "lumpy" waves, "mobbing" technicians during some intervals, such as after new malware is released, and leaving them largely idle during other hard-to-predict periods.)

    - We are assuming the technicians work forty-hour work weeks and no time is lost to meetings, training, vacations or other normal non-productive periods. That implies the center would need a minimum of (100,000/40)=2,500 technicians per week, year in and year out. If the center hires moderately skilled employees, we might expect a fully-loaded salary cost of $60,000 per technician, including direct salary, benefits and required equipment, such as computer, network connectivity, telephone access, office furniture, software, documentation and other necessities.
    - $60,000*2,500=$150,000,000/year. That's a daunting budget for these budget-constrained times.

- Anyone who has worked in technical support can tell you that it is very inefficient to walk users through a complex, multi-step technical process over the phone. If the proposed model is one where all contact is via telephone, fax or email, our initial estimate of an hour per user is truly wildly optimistic. If the proposed model envisions users bringing their systems into a local office for help or technicians making house calls, the center will need to be broadly distributed, the cyber equivalent of USDA Cooperative Extension System Office agents who are trained to diagnose and recommend specific treatments for crab grass and other agricultural or horticultural maladies.

  Due to the complexity and cost of such a system, we assume this is not the envisioned model. Rather, we expect the hope is that users will be able to get help online, then self-remediate once they have received some initial tips. In fact, however, this approach is impractical:

    - Most users will not have even the basic software tools they might need to remediate an infected system and the presence of malware may make it difficult or impossible for them to download these. This may be due to malware actively blocking access to known remediation resources or because their ISPs have disabled their access to the Internet once malware was detected on their systems.
    - Even if users have some of the more-commonly suggested tools for tackling an infection, such as an antivirus program updated with the most current antivirus definitions, running these programs may not find and remove all the malware infecting their system. This is why most computer security experts recommend users "nuke and pave" (format and reinstall) an infected system rather than trying to remove malware without rebuilding. This strategy presumes the existence of clean and trustworthy recent backups of users’ files, backups that may not exist.
    - Even if a system can be disinfected or rebuilt from scratch with clean backups, the system also needs to be patched and otherwise hardened or it may only be a matter of time before it is re-infected. Patching and hardening a system requires a careful balancing of user requirements and user risks, a balance that varies from user to user and a realm that is not easily standardized into a one-size-fits, readily-automated technique.

In many cases, it is likely that ISPs will be less enthusiastic about a centrally-operated system because this may have the effect of breaking the link between the customer and the ISP. Many ISPs are already addressing bot notification systems and these are showing beneficial effects that range from an opportunity to sell additional security services beyond simple self-remediation to reducing churn in the customer base.

IV. "General Questions on Practices To Help Prevent and Mitigate Botnet Infections" Section

(1) What existing practices are most effective in helping to identify and mitigate botnet infections? Where have these practices been effective? Please provide specific details as to why or why not.

Specific techniques used by ISPs and security vendors to identify and mitigate botnet infections are seldom publicly discussed in detail out of a desire to avoid inadvertently helping bot herders evade detection, and to protect sources and methods. For these reasons, and mindful of the fact that this response will be publicly available, we will limit our response to the first part of this question to the following basic points:

(a) We assume that we are discussing a network that is "well instrumented," e.g., one that at least routinely collects Netflow data or the equivalent for the traffic that passes over it. Note that it may not always be easy to extensively instrument one's network, particularly if:

- The network is operating at 40Gbps or 100Gbps, speeds at which monitoring may be technically challenging, or
- The as-deployed network architecture requires instrumentation at a large number of discrete locations, rather than a relatively small number of core nodes, or
- The network carries IPv6 or other advanced protocols, protocols which may not be readily measurable on some networks, or
- In many cases, the most obvious methods of detection such as DPI have been removed from the toolkit of ISPs because of privacy concerns or regulatory pressure.

(b) The easiest way to identify bots in bulk is by monitoring network traffic for the bots emissions. For example, if a botnet is used to send spam, the spam emitted by that host is the telltale indicator demonstrating the host is botted. This type of spamming may be directly recognized by Netflow processing software or the spamming may be noticed as a result of other factors, potentially including details as simple as complaints from other sites or distinctive patterns associated with the bots use of DNS.

Not all types of bot emissions are directly attributable to their actual source. For example, consider a DDoS (Distributed Denial of Service) bot emitting spoofed UDP[1] traffic: once that traffic exits the source network, it may be difficult or impossible to directly attribute its true origin. This would be less of a problem if more networks employed BCP38[2] filtering of spoofed traffic, but as noted by the MIT Spoofer project (http://spoofer.csail.mit.edu/index.php), over 20% of all autonomous systems continue to allow the emission of spoofed UDP traffic by their customers. This is a concrete example of a condition a voluntary ISP code of conduct might explicitly seek to change.

(c) Once a botted host has been identified from its emissions, one approach to identifying additional hosts that are also botted – even if they are not yet emitting spam or other problematic traffic – is to "look upstream." That is, bots tend to be controlled via commands issued from “command and control” (C&C) hosts. Botted hosts "check in" with their C&C hosts and then subsequently receive instructions from that C&C. Those instructions might direct a bot to update or extend itself by downloading new software, tell the bot to send spam, or to launch a DDoS attack at a particular targeted site. Obviously, once an ISP knows a C&C exists, communication between that C&C and one of its user’s systems is a sign that something is likely seriously amiss.

With respect to the questions, "Where have these practices been effective? Please provide specific details as to why or why not," again, recognizing that this is a sensitive question, we will limit our remarks to the following points:

(d) Networks have to ***be interested in identifying, notifying and mitigating botted customer hosts.*** If network operators are not interested in doing so, botnet identification, notification and mitigation will not occur. We have distinguished three main reasons for a potential lack of interest:

- Dealing with an infected customer is potentially expensive and, if not required, represents an avoidable expenditure that could otherwise reduce the profitability of the ISP. Valid, quantifiable research on the true costs of supporting customers in this fashion would go a long way to containing the justifiable concerns of ISPs in this area.
- While the early tradition on the Internet was that each provider had an obligation to be a "good neighbor," taking care of any abuse originating from its facilities, later Internet entrants often brought with them a different perspective; e.g., that they were a quasi-common carrier and their responsibility was simply to provide connectivity to users on a non-discriminatory basis with no obligation (or even ability) to "interfere" with how that connectivity got used. If you have the perspective of those later entrants, policing the Internet

[1] User Datagram Protocol; see https://secure.wikimedia.org/wikipedia/en/wiki/User_Datagram_Protocol

[2] RFC3704/BCP 38, "Ingress Filtering for Multihomed Networks," March 2004, http://www.ietf.org/rfc/rfc3704.txt

was the responsibility of, well, the police, even if law enforcement had no mandate or ability to operate effectively in that space.

- Some companies have been known to intentionally turn a "blind eye" to what happens on their networks, receiving a premium price from malicious customers for offering "bullet proof" service.

(e) Network-based bot detection solutions require a ***well-instrumented network***, as discussed above. If a network is not built and configured around a design that allows the network operator to monitor its own facilities, the operator will be "driving with blinders on" and will have a harder time identifying botnets established among their customers or confirming complaints they may receive. It is important that methods of network management not be dictated to the ISP but that network management principles be directed toward protecting user privacy while allowing the ISP to effectively enhance user security.

(f) Network operators need the ability to ***communicate effectively with their customers***, most typically via email. That ability may not always exist.

While many operators may routinely provide their customers with email accounts, customers may not use them, perhaps preferring an email account obtained at work or school or a free account from Gmail, Hotmail, Yahoo or some other provider. The network operator may have no idea what a specific customer's preferred email account – the one they actually use and read – might be.

Even if the network operator does know a user's preferred email address, customers may be wary of trusting or relying on "warnings" received by email since they will often have routinely received wave upon wave of phishing messages. They may not be able to distinguish a real warning from a fake one.

Alternative communication channels, such as calling the customer on the phone or sending the customer a written letter by U.S. mail, may be too expensive. These also may be routinely ignored or disregarded as potential marketing contacts by their ISP.

Ultimately, the contacts that are hardest to ignore and most credible are probably those that interfere with the surfing process, either as a captive portal or messaging on screen, such has been detailed in IETF RFC 6108[3]. A captive portal is often referred to as a “walled garden.”

As part of that process, users may receive advisory communications directly in their Web browser, in lieu of (or in addition to) their normal home page or a site they expected to visit. By demonstrating control over the user's Web experience, such messages are both hard to ignore and demonstrably credible.

(g) Network operators also need the ability to communicate and share cybersecurity data with each other. For example, if an ISP is seeing incoming spam from another network operator, the ISP receiving that spam should ideally be able to share evidence of that abuse with the originating ISP. Attempts to create data sharing arrangements of that sort may be challenged as a result of technical and policy-related concerns, including:

- A lack of usable points of contact for some providers; RFC2142[4] role accounts may not exist, or if they exist, they may not be read and acted upon.
- Some ISPs may apply consumer grade anti-malware and anti-spam filtering to abuse reporting addresses, thereby making it hard to report malware and spam with the necessary evidence.
- Different providers may employ or prefer different reporting formats. For example, should reports be exchanged in RFC5965[5] "ARF" format, RFC5070[6] “IODEF” format or is another format needed?

[3] "Comcast's Web Notification System," February 2011, http://tools.ietf.org/rfc/rfc6108.txt

[4] "Mailbox Names for Common Services, Roles and Functions," May 1997, https://www.ietf.org/rfc/rfc2142.txt

[5] "An Extensible Format for Email Feedback Reports, August 2010, https://tools.ietf.org/html/rfc5965

- When abuse reports are shared, are they shared in an un-redacted format or do privacy considerations mandate redaction of potential customer-identifying data?
- Should all abuse reports be forwarded or just a representative sampling, such as 1-in-10 or 1-in-100? If too many reports are forwarded, it may have the effect of DDoS'ing the abuse reporting address; if too few reports are forwarded, some particular abuse incidents may end up “slipping through the cracks."

Some ISPs have successfully tackled these issues, and Spamhaus lists 15 providers of such "feedback loops" at http://www.spamhaus.org/faq/answers.lasso?section=ISP%20Spam%20Issues#119 Encouraging ISPs to offer feedback loops, to have monitored abuse reporting addresses, and to avoid over-filtering those abuse reporting addresses are good examples of potential elements of a voluntary code of conduct for ISPs. Consideration should also be given to encouraging ISPs, hardware vendors and software developers to work in the IETF to develop a protocol for rapid communication of threat data for automated consumption and action. In this regard, a recent grant from the National Science Foundation is helping to support the SES-Security Event System (see http://www.ren-isac.net/ses/), an open source threat information sharing system. (See http://www.ren-isac.net/ses/ses_news.html for details.)

(h) Sometimes simply identifying an actual botted host may be difficult. This is particularly true for sites that make extensive use of NAT with PAT[7], sharing a single public IP addresses with scores or even hundreds of internal hosts. Which "needle in that haystack" is actually infected? As IPv4 exhaustion plays out and use of IPv6 or carrier grade NAT becomes common, this issue will become only more common.

(2) What preventative measures are most effective in stopping botnet infections before they happen? Where have these practices been effective? Please provide specific details as to why or why not.

Many practices are effective, or at least potentially effective, in stopping botnet infections before they occur. Some of those practices include:

(a) Regardless of the specific operating system, it is critical that systems be kept ***patched and up to date***, including the operating system itself as well as all applications such as office suites and browsers that may be installed. This also includes installed helper applications such as Adobe Acrobat Reader, Flash, Java and QuickTime. Secunia PSI/CSI/OSI or other patch management tools may be tremendously helpful in identifying installed applications in need of patches.

(b) Systems often come with many unneeded services or protocols enabled by default. This increases the system's potential "attack surface." Thus, an important step is to ***disable any unneeded services or protocols***. Put tangibly, if a user self-tests their system with a port scanning utility such as GRC's ShieldsUp (https://www.grc.com/x/ne.dll?bh0bkyd2), few if any ports should show as open.

(c) Hardware or software ***firewalls*** have historically been another staple "common sense" recommendation, and were potentially helpful in deflecting historic "scan and sploit" remote attacks.

Unfortunately, traditional firewalls are unable to protect systems against malware that is dropped when a user unknowingly visits a tainted Web page. Firewalls may also limit throughput, interfere with H.323 video conferencing or the use of native IPv6, and serve as a potential stateful[8] chokepoint for denial of service attacks. Their value is thus something of a mixed bag at this point, even though many users continue to rely on them, at least at home. Laptop or mobile device users may have relatively few options

[6] https://www.ietf.org/rfc/rfc5070.txt

[7] See https://secure.wikimedia.org/wikipedia/en/wiki/NAT

[8] See https://secure.wikimedia.org/wikipedia/en/wiki/Stateful_firewall

when travelling and connecting via third-party wireless access points, except for on-device software firewalls.

(d) Antivirus products routinely get mentioned as an important part of every user's strategy for avoiding malware infections; however, increasingly, the pace of malware production and release is outstripping the ability of some antivirus companies to keep up. As a result, even users who are running a leading commercial antivirus product with up-to-the-minute antivirus definitions may still be vulnerable to infection by malware. Users have been educated to believe that antivirus is the total solution; this is no longer the case.

(e) Users need to use ***strong – that is, long and complex – passwords*** for their systems and online accounts, and need to not allow those passwords to be sniffed or phished. When supported, two-factor authentication, such as the use of a hardware cryptographic fob in addition to just a password, or two channel authentication, such as use of a code sent to a user's smart phone in addition to just a password, can dramatically reduce the ability of bots to take over user accounts.

(f) The principle of "least privilege" means that if users do not need particular privileges, they should avoid using them when performing routine tasks. In this context, users can minimize the potential impact of any malware they encounter by ***not running as a privileged user*** ("admin," "root," or the equivalent) ***except when necessary*** for software installation or other maintenance.

(g) Javascript is pivotal to many highly-popular interactive sites. However, Javascript can also serve as a common path for dropping malware on user systems. Some users have thus taken to either ***blocking Javascript*** outright, even if that means forgoing some hugely popular websites, or using a tool such as NoScript for Firefox to block Javascript when a site that is not trusted attempts to use it for potentially malicious purposes.

(h) Many useful sites are underwritten by revenue from online advertising. However, recognizing that advertising is not essential to most users and that malware can be unintentionally delivered via advertising sites, some users who are uninterested in seeing these ads use tools such as AdBlock Plus for Firefox to ***block content from advertising sites***. In doing so, those users reduce the likelihood they will be a victim of "malvertising," malicious software delivered via advertising sites.

(i) Avoid using peer-to-peer applications or trusting files downloaded from P2P sites.

(j) Another approach is for ***ISPs to deny malware authors access to reliable DNS*** (domain name system). Currently, most sites faithfully resolve malicious domain names just as they would the domain name of legitimate sites. Paul Vixie of the ISC[9], author and proponent of RPZ ("Response Policy Zones"), urges ISPs and users to deny miscreants the ability to rely on DNS. Specifically, if ISPs blackhole (filter) known malicious DNS resource records, cyber criminals will no longer be able to use DNS to resolve malware dropping sites, command and control hosts, phishing sites, and other malicious locations.

(k) Consideration should be given to "locking" users to a specific DNS that is known to be reliable. Most users do not care about the source of their DNS and malware authors take advantage of this. By restricting the end user on an opt–out basis to a reliable DNS, it is going to further restrict the malware authors from exploiting the DNS. We also congratulate the federal authorities on their successful take down of the DNS Changer malware gang via "Operation Ghost Click" (see www.fbi.gov/news/stories/2011/november/malware_110911/malware_110911).

(l) Use of a more modern version of Windows will reduce your risk of being botted relative to use of an earlier version of Windows. See, for example, "Security Intelligence Report, Volume 11," covering the period from January to June 2011, at http://www.microsoft.com/security/sir/default.aspx.

[9] Internet Systems Consortium (https://www.isc.org/).

> That report notes that the infection rate by operating system and service pack continues to improve (decrease) as users move from Windows XP, to Windows Vista, to Windows 7. (See PDF page 82 or print page number 58.)

Users and sites have tried all those approaches, and many more, in an effort to preemptively block botnets. When employed conscientiously, they typically result in a material reduction in hosts being botted. When they are employed less conscientiously, less protection is delivered and more compromised hosts are observed.

A good starting point for users interested in how they can harden their systems can be found in The Security Configuration Guide at the National Security Agency Central Security Service site. See http://www.nsa.gov/ia/mitigation_guidance/security_configuration_guides/operating_systems.shtml

(3) Are there benefits to developing and standardizing these practices for companies and consumers through some kind of code of conduct or otherwise? If so, why and how? If not, why not?

To answer this question we must consider what happens if a voluntary code of conduct is ***not*** developed and standardized. If such a code of conduct is not developed and standardized, one potential result might be mandatory federal regulation. As the lesser of two evils, most ISPs would agree that a voluntary code of conduct is preferable to mandatory federal regulation, particularly if those regulations represent an unfunded mandate, and since mandatory federal regulations generally are unable to keep pace with the rapidly evolving threats that ISPs constantly face.

Arguing in the alternative, given that we have seen consistent gridlock at the federal level when it comes to passing cybersecurity legislation, we must also recognize the possibility that mandatory federal cybersecurity regulations may be simply too contentious to be enacted at this time. In such a case, imperfect though a voluntary code of conduct might be, it might be the best option for clearly defining minimum community expectations for cybersecurity.

In either case, a voluntary code of conduct will likely be embraced most strongly by those ISPs that already proactively take all feasible steps to deal with botted customers. Those who are ambivalent, or who actively conspire with cyber criminals, are unlikely to change their behavior as a result of the creation of a voluntary code of conduct, absent economic incentives.

We also need to recognize that some ISPs may claim they intend to follow a voluntary code of conduct while failing to actually do so. Some thought should be given to how agencies might deal with entities that claim to adhere to a voluntary level of performance while flouting the terms of that code in practice.

(4) Please identify existing practices that could be implemented more broadly to help prevent and mitigate botnet infections.

Please see our responses to questions (1) and (2) above.

(5) What existing mechanisms could be effective in sharing information about botnets that would help prevent, detect, and mitigate botnet infections?

We distinguish between two approaches to "information sharing about botnets" as:

- General best common practices (BCPs) for preventing, detecting and mitigating botnet infections, and
- Specific operational details about observed botted hosts, botnet command and control hosts and their name servers, sources of botnet infections, and related specific data-driven "actionable intelligence."

MAAWG, and other industry organizations, routinely share BCP information. This is an effort to help our members and third parties cope with the problems that botnets can cause.

Specific data-driven actionable intelligence about botnets is also broadly shared within the community. We already mentioned that some ISPs have created feedback loops designed to share information about botted hosts that emit spam [see IV. (1) (g), above]. This is an excellent example of distributing current intelligence. Various DNS-based blocklists, such as those offered by Spamhaus[10] or SURBL[11], are another example. If an ISP has IPs or customer domains on those blocklists, they probably have botted customers. Data-driven operational intelligence can also form the core of a proprietary for-fee security service from a commercial entity or may serve as the entire *raison d'être* for the purchase of some types of network security appliances.

In other cases, actionable intelligence about botnets may be shared at no cost but only on a need-to-know basis. The goal in this case typically is to protect the sources and methods used to develop that operational intelligence or to avoid retribution from disgruntled bot herders.

(6) What new and existing data can ISPs and other network defense players share to improve botnet mitigation and situational awareness? What are the roadblocks to sharing this data?

ISPs and other community participants can share:

> ***(a) IP addresses*** with timestamps of botted hosts observed to be spamming or otherwise misbehaving (see for example http://www.spamcop.net/w3m?action=inprogress).
>
> ***(b) Spamvertised domain names*** or other points of contact where replies are being directed, such as phone numbers or throw away email addresses; see, for example, http://www.spamcop.net/w3m?action=inprogress;type=www,
>
> ***(c) Botnet command and control host information*** (see http://mtc.sri.com/live_data/cc_servers/).
>
> ***(d) Name servers*** used by the above.
>
> ***(e)*** Information about ISPs that may be routing botnet-related infrastructure (e.g., ***ASNs*** that are routing botnet-related netblocks).
>
> ***(f)*** Attribution (owner/user) information associated with botnet-related assets, including ***IP WHOIS*** and ***domain WHOIS*** information, and any shell companies or fake identities used to purchase domains names, hosting services, or other infrastructure used by bot herders or bot users.
>
> ***(g)*** Details of affiliate programs associated with bot-spammed email (e.g., ***affiliate codes*** associated with spam).
>
> ***(h) Malware samples (or hashes)*** associated with botnet infections: analyses/reverse engineering write-ups associated with such malware.

(7) Upon discovering that a consumer's computer or device is likely infected by a botnet, should an ISP or other private entity be encouraged to contact the consumer to offer online support services for the prevention and mitigation of botnets? If so, how could support services be made available? If not, why not?

Yes, we believe the consumer ***should be notified*** whenever possible. However the details of the support services, if any, that may be appropriate beyond this notification will vary depending on the circumstances of the observed bot.

As previously discussed, in many cases botted hosts are not easily disinfected. It is ***not*** just a matter of telling users, "Download the antivirus product we have provided for your use and run that software to clean up your system" any more than a user should be expected to self-treat tuberculosis or cancer. It would be extremely valuable if there was a "magic bullet" that could cure all malware (or all human diseases) but we have yet to discover such a magic potion.

[10] See http://www.spamhaus.org/

[11] See http://www.surbl.org/

In the case of a malware-infected system, many things can prevent a user “self-help” approach:

(a) *The malware may interfere* with the user downloading or executing the recommended antivirus product.

(b) If the user is able to download and run the antivirus product, it may not *detect* the malware that's infecting the host.

(c) If it does detect the malware, it may not be able to safely, or completely, *remove* it.

(d) If the user does succeed in removing the malware, the system may still be infected with *other malware*.

(e) The malware just removed may quickly be *reinstalled* unless the system's underlying vulnerabilities are also corrected.

Bottom line: Online self-remediation is simply *not* sufficient to cure most botted hosts.

(8) What should customer support in this context look like (e.g., web information, web chat, telephone support, remote access assistance, sending a technician, etc.) and why?

Most commercial malware remediation services ("rent-a-technical-support-person" outfits) offer three basic models:

- Help at the user's site–they make a "house call" to the user’s home or business. This is typically the most expensive option.
- Help at the service provider's location–you bring your system to a store front location. This is the least convenient option from a user's point of view, at least for desktop systems.
- They work with you online–this option presumes that your ISP has not cut off your connection and that malware is not actively interfering with remediation efforts. This is typically the least expensive option.

We assume that most users prefer the least expensive and easiest option for their particular circumstances. If a consumer is paying out-of-pocket for remediation service, how much might they be willing to pay? The rational upper limit probably is:

> The cost of a new system
> + The cost of transferring critical applications and user files from the old system or backups
> –The expected salvage (resale) value of their old system
> =The rational ceiling on potential remediation expenditures

Users may also perceive an intangible value to buying a new system that may be faster, more reliable or have more features than the old system they are replacing. Perfectly serviceable new laptops are available for $500 or less and new desktops are less expensive. If a system is severely infected or several years old, the user may consider replacing it rather than spending hundreds of dollars either attempting to clean it or executing a nuke-and-pave strategy.

But the preceding is from an end-user's point of view. From the ISP’s perspective, the form that customer support should take depends on:

- The scale at which the ISP needs to provide support (one user? a hundred users? ten thousand users?)
- The ISP’s goals in providing that support
- The technical expertise of their customers
- The specific characteristics of the malware infecting their customers’ systems

If a small local ISP only needs to help a handful of botted customers a day, having an otherwise not busy technician informally make a house call may provide the best customer experience. However, if the provider has thousands or tens of thousands of customers who need help each day, and those customers are spread across a broad geographic region, it may not be practicable to dispatch a technician to fix their infections on site. More accurately, it would be prohibitively expensive to do so.

Conservatively, even the briefest of house calls costs a minimum of $75, considering costs such as staff salary and, benefits, providing a vehicle and fuel, staff hardware and software tools, insurance against errors and omissions, and other factors. Given today's highly competitive marketplace and slim margins, the expense associated with a single house call would likely wipe out any profit from that customer for months, if not years, to come.

Passive delivery of remediation information to consumers, e.g., via Web pages, presumes an availability of tools and expertise that will only rarely be present. Without having those tools and the necessary specialized expertise, users will likely be overwhelmed attempting to self-remediate and may actually make their circumstances worse.

Synchronous support via the telephone often is prolonged and takes more time, as users need to locate and relay what they see on their screen. This is particularly true with non-technical users or users anxious about the possibility they will "break something."

Attempting to short circuit that process by installing a remote access client, allowing a technician to temporarily take control of the user's system for diagnosis and remediation purposes, raises questions of user privacy and trust. In most cases, remote access clients give the off-site technician complete access to the user's system, including all the potentially sensitive information that may be stored on it such as personal email, documents and photographs, saved passwords, past tax returns, and other private data. How is the user to know their sensitive information and irreplaceable content will not be leaked or accidentally lost? How is the user to know the remote access capability has been disabled or uninstalled when it is no longer needed? This service currently is widely offered, and should always be offered, with the appropriate warnings about privacy and a full opt-in from the user.

(9) Describe scalable measures parties have taken against botnets. Which scalable measures have the most impact in combating botnets? What evidence is available or necessary to measure the impact against botnets? What are the challenges of undertaking such measures?

We might best begin by considering what makes a potential anti-bot measure scale poorly:

- A human being is involved somewhere along the line (e.g., either the user or an ISP customer service technician needs to take some action). The ISP needs to leverage automation rather than rely on slow and error-prone human intervention.
- A systematic understanding of the underlying phenomena is lacking and, as a result, structural weaknesses of eventual exploitation have not been identified. For example, imagine that as botnets are leveraging fast flux[12] hosts, defenders are attempting to identify and take down participating hosts one at a time rather than going after the fast flux name servers or domain name registration.
- Poor scalability is often associated with a solution that needs to be custom crafted for each case; scalability suffers if there is no general-form solution that is broadly applicable.
- State needs to be maintained and "remembered" somewhere (access control lists in a firewall, router or other network appliance; whitelists or blacklists of unwanted IP addresses or unwanted domain names; etc.)
- State needs to be replicated and updated at many points, rather than being maintained at a centralized location (how will you replicate and distribute that state to all the points that need a current copy of it?)
- State is not able to be cached so that the same query (and the same answer) might get asked and answered repeatedly.
- There is no strategy for "garbage collection" or cleaning saved state information that is no longer timely, resulting in accumulated state becoming diluted by irrelevant legacy information that should have been purged.
- The strategy has a tangible cost (e.g., mailing a CD to each infected user scales badly because of the costs of creating and delivering those CDs).

[12] A technique to hide phishing and malware delivery sites. See https://secure.wikimedia.org/wikipedia/en/wiki/Fast_flux

Given those considerations, one can quickly see why truly scalable anti-botnet measures are usually either:

- Network-based solutions, such as DNS-based blocklists. For example, Spamhaus currently protects 1,438,371,000 user mailboxes as of October 18, 2011 (http://www.spamhaus.org/organization/index.lasso), or
- Automated tools, such as the Microsoft Malicious Software Removal Tool that runs each month as part of Microsoft Update. Microsoft notes that the MSRT was downloaded and executed 3.2 billion times in the first half of 2010 alone. (See Microsoft Security Intelligence Report, Volume 10, January-December 2010, PDF page 72, as described below).

A concrete example of a less-scalable anti-botnet measure is the FBI's Coreflood botnet takedown (see http://www.fbi.gov/news/stories/2011/april/botnet_041411). While that effort was and is appreciated, and had a material affect on botnets, it required excessive effort by federal agents. If you are not familiar with that takedown, see, for example:

- "FBI vs. Coreflood botnet: round one goes to the Feds," http://arstechnica.com/tech-policy/news/2011/04/fbi-vs-coreflood-botnet-round-one-goes-to-the-feds.ars
- "FBI Begins Disinfecting Coreflood from User PCs," http://securitywatch.eweek.com/botnets/fbi_begins_disinfecting_coreflood_from_user_pcs.html

Why was that takedown less scalable?

- Extensive legal pleadings were required.
- An attempt was made to allow users to "opt out" of having Coreflood removed – an option that few if any would knowingly exercise – and to explicitly solicit user consent to clean the infection from their system.
- Anti-bot measures were only applied to infected systems located within the United States.
- Ultimately many instances of Coreflood were deleted by the MSRT rather than by the direct actions of the FBI.

When it comes to measuring the effectiveness of these approaches, several metrics are available. For example, consider the Composite Block List (the "CBL", see http://cbl.abuseat.org/), the botnet-focused component of the Spamhaus blocklist. The CBL website offers a variety of metrics, including:

- The distribution of botted hosts per domain, per country, per ASN, etc.
- A breakdown showing the botnets that CBL is routinely seeing
- A pointer to a comparison of various blacklists, see http://www.sdsc.edu/~jeff/spam/cbc.html
- Many other fascinating statistics

Microsoft also routinely shares information about the impact of their Malicious Software Removal Tool as part of their periodic Security Intelligence Report (SIR); see http://www.microsoft.com/security/sir/default.aspx. (Note: both a brief summary and a far more detailed report are provided at that site. Be sure to see the detailed report for detailed MSRT-derived metrics.)

V. The "Effective Practices for Identifying Botnets" Section

(10) When identifying botnets, how can those engaged in voluntary efforts use methods, processes and tools that maintain the privacy of consumers' personally identifiable information?

We assume that the departments issuing this RFI did not use the term "personally identifiable information" casually, but rather as a term of art. We therefore need to begin by defining what we consider to be "personally identifiable information" (PII). Sensitive categories of PII are normally considered to include data such as:

- Credit or debit card numbers, particularly if the card holder's name, billing address, card expiration date and CVV codes are included
- Government assigned identifiers, such as SSNs, drivers license numbers, passport numbers, etc.
- Health-related information; e.g., medical records covered by HIPAA
- Financially-related information; e.g., bank records and other information covered by GLBA
- Educational records; e.g., FERPA-covered data
- Biometric or genetic data about a person or other physical and personal data

The sensitivity of these categories increases when combined with other data or user attributes such as the user's name, date of birth, likeness, height and weight, hair and eye color, gender, race, marital status, disability status, veterans status, home address, telephone number, preferred email address and other information.

Most types of sensitive PII are not routinely seen when exchanging information on botnets – or at least not in conjunction with consumers or other innocent third parties. The one potential exception may be user IP addresses. Some jurisdictions hold that an IP address is potentially PII, while other jurisdictions have determined that it is not. (Note: A fundamental part of anti-botnet work does often involve accumulating details about known or suspected bot operators or supporting organizations, including affirmative efforts at preventing those cyber criminals from remaining anonymous and unattributed. For example, see the Spamhaus Register of Known Spam Operations (http://www.spamhaus.org/rokso/).

In fact, one of the reasons ISPs are often viewed as pivotal in the fight against botnets is that average users enjoy a high degree of *de facto* anonymity online, unless they engage in activities that undercut that status. For example, most users will automatically receive a dynamic IP address when connecting to their ISP and will use that dynamic IP address while working online. After a period of time that varies by ISP from mere minutes to multiple days or longer, the "lease" for that IP address will either be renewed or relinquished. Customers commonly receive a different IP address the next time they connect, which is why the addresses are known as "dynamic," rather than "static.”

While the user's ISP can normally map network activity associated with an IP address at a given point in time to a customer, only authorized employees of the ISP with privileged access to the necessary network logs and business records will generally be able to do so. ***External parties*** will generally NOT be able to map IPs to names, although there are exceptions. These exceptions arise from users’ voluntary behavior and decisions. For example, a third party could potentially map an individual's identity to their IP address when a user logs into a website account that is tied to the user’s actual identity, perhaps as a result of the user making a purchase there with a credit card. Once that happens, the website operator has the ability to map that logged-in user's identity to their IP address, at least until such time as the user is given a new IP address. If the user remains logged in or cookies are being used, the website may be able to continue to know the user's identity even as the user moves from one IP address to another.

Bottom line: Consumers’ personally identifiable information is typically ***not*** at risk as a result of botnet identification-related efforts.

(11) How can organizations best avoid “false positives” in the detection of botnets (i.e., detection of behavior that seems to be a botnet or malware-related, but is not)?

False positives are rare when sound behavioral approaches are used to locally identify bots based on network traffic. This is because normal customer network traffic does not look anything like that emitted from a botted host. However, it is possible to conceive of some situations with potential false positives, particularly in conjunction with external botnet notifications. For example:

- We must acknowledge the possibility of malicious and unfounded botnet-related reports, perhaps made in an attempt to "punish" security researchers or others who are interfering with the conduct of online cyber crime. These reports are easily discredited by the ISP if internal Netflow data is available and the externally reported incident(s) cannot be corroborated.

- Sometimes there may be accidental confusion over time zones when bot detections are made or clocks may not be well anchored to trustworthy NTP sources.

  When that happens and a dynamic address is involved, the mapping of an IP address to a user may go awry, pointing to an innocent user rather than an actually botted one. (Remember that a dynamic address may be used by one user at a given point in time, then by some other user shortly after; distinguishing the two users requires accurate time stamps). Careful attention to time zone issues, attention to NTP synchronization, and work to confirm external reports using internal Netflow data can mitigate these potential false positives.

- BGP route injection is another approach that may potentially obfuscate the true source of bot traffic. If an abuser with a complicit ISP announces a more specific route for an already-in-use netblock, and then spams (or engages in other undesirable behavior) from that hijacked address space, the actual authorized user of that address space might end up getting "blamed" for the abuse even though that abusive network traffic may never have been anywhere near the actual authorized user's network. (For a discussion of this, see "Route Injection and Spam," (http://pages.uoregon.edu/joe/maawg8/maawg8.pdf). Careful monitoring of wide area BGP routing data and verification of external abuse reports against internal Netflow data can largely mitigate this risk.
- A malware researcher might intentionally download botnet-creating malware in a controlled way for analysis purposes. If the network monitoring equipment in use does not know this is intentional behavior by a security researcher and not an accidental infection of an average user, it may flag that researcher's activity as a possible new bot infection.
- Primitive (and not-recommended) "botnet detection" heuristics can create situations such as mistakenly flagging any user connection to an IRC (Internet Relay Chat) server as a potential botnet command and control connection. While IRC is an admittedly popular botnet command and control technology, IRC can be, and is, used for legitimate purposes by hundreds of thousands of users worldwide every day.
- Naive reporters may not understand that UDP traffic, unlike TCP traffic, can be spoofed if the emitting network does not follow BCP38[13]-recommended filtering practices. Thus, from time to time, there may be reports complaining about spoofed UDP traffic that "you" emitted.
- Non-technical reporters may also fail to understand that email addresses and email message "Received:" headers (except for the handoff host) are also usually easily spoofable. This may result in potential false positives if the headers associated with bot-emitted spam are not correctly interpreted.

Finally, this question might equally well ask, "When detecting bots, is ***collateral damage*** a possibility?"

Collateral damage, or inadvertently hurting a user that is not botted when dealing with a user that is botted, is a distinct possibility, particularly if users connect multiple systems via a single shared IP address (e.g., by using NAT with PAT). Imagine a situation where an ISP detects bot traffic from its customer's IP and reacts by blocking that IP. If that IP is shared by multiple downstream customer systems, efforts to block bot traffic from one of those downstream customer systems may unavoidably end up blocking legitimate traffic from other downstream customer systems sharing that same IP address.

(12) To date, many efforts have focused on the role of ISPs in detecting and notifying consumers about botnets. It has been suggested that other entities beyond ISPs (such as operating system vendors, search engines, security software vendors, etc.) can participate in anti-botnet related efforts. Should voluntary efforts focus only on ISPs? If not, why not? If so, why and who else should participate in this role?

ISPs play a potentially pivotal role in detecting bots because they may be the only entity able to map the IP of a botted host to the identity of an actual customer, and thus may be the only ones able to directly contact a botted customer about their infected status. However, others also have roles to play in combating botnets. Efforts must **not** be limited just to ISPs.

[13] RFC3704/BCP 38, "Ingress Filtering for Multihomed Networks," March 2004, http://www.ietf.org/rfc/rfc3704.txt

Combating the menace of bots can best be described as a team effort: ISPs have a leading role in the detection of botnets and notification of infected users. Tools vendors have a leading role to play in the removal of malware from infected hosts. Operating system vendors need to ensure that opportunities to compromise operating systems are minimized. Application vendors need to ensure their applications are both secure and easily patched. This is currently a problem as the multitude of patching methods is confusing and may be the subject of spoofing attempts.

We have previously mentioned the good work that Microsoft has done with its Malicious Software Removal Tool. Clearly their efforts to directly remove malicious software from user systems with MSRT are tremendously helpful in the fight against bots, as is their program of civil litigation against botnet operators.

Beyond ISPs and Microsoft, one needs to recognize that botnets operate as part of a highly specialized ecosystem. If bots are to be commercially profitable for their operators, a variety of conditions need to be satisfied; disrupting one or more of those conditions may be sufficient to make botnet operations financially unprofitable.

For example, we know that while bots may drive some traffic to spamvertised hosts, additional financially-important traffic comes from search engine listings. If search engines delist (or at least "down rank" or "deprioritize") pages from websites that are being spamvertised by bots (whether that spamvertising is being done by email, Web spamming tools or other means), that sort of technical "penalty" can be devastating to the commercial viability of bot-promoted enterprises.

Similarly, hosting companies also have a role to potentially join in: managing botnets, particularly large botnets, requires high capacity/stable servers and substantial bandwidth. That requirement is normally satisfied by bot herders purchasing "bullet proof hosting" from complicit, or at least "intentionally disinterested," hosting companies. If hosting companies refused to provide bot herders with the facilities they need for botnet administration purposes, large botnets would be difficult or impossible for bot herders to administer. While some hosting companies may be innocent or unknowing victims of botnet operators, hosting companies that knowingly tolerate botnet management servers in their facilities or on their networks need to be held accountable for that business decision. In particular, hosting companies, and any resellers or agents thereof, should be required to have and follow strict 100% "know your customer" policies. Specifically, this means that it should be impossible for a customer to acquire hosting for immediate use without closed-loop confirmation of the customer's identity.

Registrars and registries also have a role. If the dedicated domain names used for botnet command and control hosts can be disabled or put on a "HOLD" status – thereby rendering those domain names unusable – bot herders will be hard pressed to maintain control over their botnets using traditional command and control approaches. We also need to ensure that all registries move to a so-called "thick registry model" so that the registry can identify not just ***one*** botnet-related domain, but the ***complete*** set of botnet-related domains that may have been registered to a given bot herder with the same (or closely related) registration details.

We also note that miscreants routinely abuse privacy and proxy registration services. Given the levels of abuse associated with such services, we believe that private/proxy registration services should no longer be allowed for commercial enterprises. Customers deserve the right to know the identity of the company they may be trading with.

We have no objection to private/proxy registrations continuing to be available for non-commercial registrants, provided they are not demonstrated to be engaging in commercial activity. If they are, the private/proxy registration for that putative "non-commercial" entity should immediately be de-cloaked.

It is important to understand that the operation of botnets is primarily aimed at generating revenue for criminal operators. Interdicting the flow of money is likely to have the greatest effect in their demise since without financial rewards there is little incentive for many botnet operators to take part in this crime.

Finally, we would like to see the creation, promotion, sale, distribution or use of bot software be made a separately chargeable felony offense, in and of its own right.

VI. The "Reviewing Effectiveness of Consumer Notification" Section of the RFI

(13) What baselines are available to understand the spread and negative impact of botnets and related malware? How can it be determined if practices to curb botnet infections are making a difference?

When it comes to getting a baseline and understanding if efforts to curb botnets are making a difference, look at the statistics available from the Composite Block List, as discussed elsewhere in this response. You need only look at http://cbl.abuseat.org/domain.html to see examples of ISPs that are totally overrun with botnets. The cleanest ISPs, and relatively small ISPs, do not show up on that page because they have few if any botted hosts.

In most cases, if anti-botnet measures are effective, infection levels should be less than 0.01% (e.g., for an ISP with a million customers, 100 or fewer customer hosts will be botted at any point in time). Infection rates that run a full percent or higher (e.g., 10,000 botted hosts per million customers) are typically indicative of ineffective anti-bot measures, or a lack of anti-bot measures.

(14) What means of notification would be most effective from an end-user perspective?

Please see our response at section IV. (1) (f) above discussing approaches for notifying end-users.

(15) Should notices, and/or the process by which they are delivered, be standardized? If so, by whom? Will this assist in ensuring end-user trust of the notification? Will it prevent fraudulent notifications?

Bot notices to users need to convey basic information about what has been detected. Standardization may be a side effect of automating notification processes, but will not necessarily convey material benefits when notifications are being made to an end user. When notifications are being made in bulk, from one provider to another, standardization may help the receiving provider automate the ingestion and processing of those notices. See our previous discussion of this point at IV. (1) (g) above.

With respect to fraudulent notifications, standardization will potentially ***facilitate*** rather than ***deter*** fraudulent notification. Consider examples from the physical world: letterhead stationary can be readily duplicated and may lend a false air of authenticity to a bogus communication. A police uniform, along with a badge ordered over the Internet, may allow a non-law enforcement person to convincingly "pretend" to be a cop, one reason why impersonating a police office is usually an offense that receives harsh treatment.

Ultimately fraudulent notifications can only be detected and ruled out by reliance on, or corroboration with, locally collected Netflow data or comparable network data sources of trusted provenance. If the local network is not well instrumented, providers may want to formally or informally track the accuracy of reports they receive from third parties. That is:

- "Uh oh, we just got another set of reports from Bob in Utah. His previous reports have always proven to be pretty accurate, so we'd better take these seriously as well," or
- "Hmm. Another report from that guy in New Hampshire who thinks our name server is somehow 'attacking' his firewall when he's trying to visit one of his favorite websites. We've given up trying to explain what's going on to him, haven't we?"

(16) For those companies that currently offer mitigation services, how do different pricing strategies affect consumer response? Are free services generally effective in both cleaning computers and preventing re-infection? Are fee-based services more attractive to certain customer segments?

We discussed pricing strategies in IV. (8) above.

Adding to this analysis in response to your question here about ***free vs. for-fee services***, free services are often not really "free." The user may need to devote substantial time and effort ("sweat equity") toward cleaning their system, and those efforts, even if diligent and long in duration, may not assure success. While working on cleaning their systems with "free" resources, users will often have to miss work, or forego spending leisure time with their family and friends. These are real costs, albeit costs that may be paid in kind rather than in cash.

For-fee services, on the other hand, may allow a specialist to just "do their magic." Like going to the dentist for a root canal, you know that it may be expensive and unpleasant, but sometimes you just need to do what you need to do (assuming you can afford to pay for the work you need).

The downside of bringing your system to a specialist, or a specialist to your system, is that you potentially lose some privacy and need to trust the technician to do only what they are supposed to do, and to do it well. An untrustworthy computer repair technician who has full access to your system could wreck considerable havoc.

(17) What impact would a consumer resource center, such as one of those described above, have on value-added security services? Could offers for value-added services be included in a notification? If not, why not? If so, why and how? Also, how can fraudulent offers be prevented in this context?

Many ISPs already have established referral relationships with third party for-fee remediation services and may routinely refer users who need or want extra help to that partner. The consumer resource center mentioned previously, if created, is unlikely to replace or significantly undercut existing sites of that nature or the business that currently flows to the private sector for-fee security services. However, unless carefully managed it will break the link between ISP and user which helps to foster user trust and maintain one of the most significant benefits to the ISP which is the reduction of user churn.

Fraudulent offers can best be minimized by listing links to legitimate sources of computer security assistance on the ISP's website and referring the user to that site, or by creating a licensing process with approved remediation providers listed in a verifiable registry.

(18) Once a botnet infection has been identified and the end-user does not respond to notification or follow up on mitigating measures, what other steps should the private sector consider? What type of consent should the provider obtain from the end-user? Who should be responsible for considering and determining further steps?

Ultimately, the ISP has both the authority and the responsibility to deal with its own infected customers.

If infected customers cannot be reached, or cannot or will not deal with their problems, the ISP ***will*** need to take appropriate unilateral measures to limit the impact of the botted customer's system to both the ISP itself and to other users of the Internet. ISPs that fail to do so quickly become "damaged goods:" their IP address space gets listed on blocklists, legitimate customers develop deliverability problems, and the ISP acquires a bad reputation. Once that happens, it does not take long for good customers to flee and a deluge of new bad customers to arrive.

Thus ISPs must take appropriate actions to deal with a botted customer, even if the customer cannot be reached or is non-responsive. This typically means putting the customer in a so-called "walled garden" where the customer's system cannot emit or accept problematic traffic. In extreme cases, or if the ISP has not yet deployed a walled garden, the ISP may end up having to completely terminate the customer's network access. This possibility will typically be included in the ISP's terms of service, and customers will have agreed to this condition at the time they signed up.

We note that completely terminating a customer's access is generally something that ISPs and customers both want to avoid. Terminated customers represent lost revenue and completely terminating a customer's broadband access may also impact other bundled services, such as cable TV or Voice over IP (VoIP) service.

If VoIP service is interrupted during an emergency period – whether that emergency is a natural disaster, man-made catastrophe, or a medical crisis – there exists the potential for serious injury to persons and property or even loss of life. Most ISPs will attempt to limit their liability in these circumstances by either exempting identifiable VoIP-related traffic from broad traffic blocks or including provisions in their terms of service prohibiting customers from exclusively relying on VoIP service for emergency-related purposes.

With respect to the question "What type of consent should the provider obtain from the end-user?" recall that most ISPs have customers agree to terms of service giving the ISP the right to manage their connection when security issues or other circumstances may require. Even with this clause in place, most ISPs will make an affirmative effort to contact the user before blocking their connection. Unfortunately, if a botted user remains persistently infected or incommunicado with the ISP, it may not be possible for the provider to communicate with the user. It is also certainly conceivable that even if the user is contacted but objects to being disconnected, an ISP might feel they have no viable alternative.

Concerning your question, "Who should be responsible for considering and determining further steps?", we believe that ideally customers should take responsibility for self-monitoring and self-securing their systems. If they fail to adequately do so, ISPs then have a responsibility to act. If the ISP in turn fails, the ISP's upstream network service provider(s), if any, may need to intervene.

However, if the customer, the ISP and the NSP (Network Service Provider) are all unable to take care of bot-related issues that arise, civil litigation, regulatory action, or criminal law enforcement involvement may be necessary. Third parties, such as blocklist operators, may also unilaterally decide to take action against botted hosts at any time in this process, notwithstanding other steps that may be undertaken.

(19) Are private entities declining to act to prevent or mitigate botnets because of concerns that, for example, they may be liable to customers who are not notified? If so, how can those concerns be addressed?

Liability for failure to notify a customer would imply that an affirmative "duty to notify" exists as a result of statutory provisions, enabling regulations or following well-settled common law precedent. We are not aware of any such existing obligation pertaining to notification of botted ISP customers. If a duty to notify ***did*** exist, a reasonably diligent "best effort" attempt to notify, even if executed imperfectly, would likely result in less liability for an ISP than if that ISP made no attempt to notify botted customers at all.

We would thus always urge ISPs to attempt to notify their botted customers whenever possible, even if they are unable to successfully do so for all types of malware or for every potentially infected customer. If the departments authoring this RFI believe that some private parties do perceive a potential liability for their anti-bot action, we encourage the departments to consider recommending a statutory amendment that would provide a safe harbor for good Samaritans who may be making a best effort attempt to notify botted users.

(20) Countries such as Japan, Germany, and Australia have developed various best practices, codes of conduct, and mitigation techniques to help consumers. Have these efforts been effective? What lessons can be learned from these and related efforts?

To empirically answer this, we look at country-level statistics for these nations from the Composite Block List relative to the United States, plus three common comparators, e.g., Canada, the UK and New Zealand (see http://cbl.abuseat.org/country.html):

DE:	13th place	143,066 listed IPs	0.129% infection rate
AU:	58th place	14,009 listed IPs	0.026% infection rate
JP:	76th place	6,257 listed IPs	0.003% infection rate
US:	20th place	100,980 listed IPs	0.005% infection rate
UK:	35th place	40,731 listed IPs	0.070% infection rate
CA:	72nd place	7,652 listed IPs	0.009% infection rate
NZ:	90th place	3,866 listed IPs	0.047% infection rate

In reviewing this table, note that top rankings (e.g., 1st place) – or larger number of listed IPs – are bad and lower infection rates are good.

When we look at that data, we note:

(a) The performance of code of conduct countries is distinctly non-uniform: There is a huge degree of diversity in the rate of infection among the three "code of conduct" countries mentioned. For example, ***Germany has 43 times the rate of infected IPs that Japan does.*** Of course, it is hard to know why such a dramatic discrepancy exists between those two countries, given that both apparently promote best practices, codes of conduct, and mitigation measures. Perhaps Japanese ISPs are less targeted by malware than German ISPs? Obviously, Japan uses Eastern character sets – perhaps malware has not been localized to use Japanese Kanji?

That hypothesis is challenged if we note the world's two most heavily botted countries, India and Vietnam (collectively accounting for over a quarter of all botted hosts worldwide), both make wide use of non-roman character sets, just as Japan does:

IN:	1st place	1,383,745 listed IPs	3.734% infection rate
VN:	2nd place	593,169 listed IPs	3.988% infection rate

It is improbable that bots would be localized for India or Vietnam but not for Japan. We can only assume that Japanese ISPs, or their customers, are culturally and socially more inclined to diligently seek out and eradicate bots than their counterparts in India or Vietnam.

(b) The United States already outperforms two of the three "code of conduct" countries by significant factors: That is, the United States actually has an infection rate (0.005%) over 5X LOWER than Australia (0.026%) and nearly 26X LOWER than Germany (0.129%).

(c) Disregarding infection rates, and just looking at raw bot counts, the U.S. still outperforms Germany: Recall that Germany had 143,066 botted IPs while the United States had only 100,980. If a site that is being hit by a botnet only cares about actual bot volumes, and not bot rates of infection, the United States is still less of a potential problem to that target than Germany.

(d) Comparing the code of conduct countries against Canada, the UK and New Zealand, there does not appear to be a meaningful or substantive difference in infection rates. If we truly want to find comparators particularly worthy of emulation, we might look at the Scandinavian countries as a group – this five-country region outperforms both Australia and Germany.

SE:	86th place	4,357 listed IPs	0.020% infection rate
NO:	111th place	1,550 listed IPs	0.015% infection rate
DK:	123rd place	1,039 listed IPs	0.008% infection rate
FI:	129th place	815 listed IPs	0.007% infection rate
IS:	154th place	172 listed IPs	0.023% infection rate

(21) Are there best practices in place, or proposed practices, to measure the effectiveness of notice and educational messages to consumers on botnet infection and remediation?

(a) It is hard to teach users everything they need to know about cybersecurity – even just ***general*** cybersecurity – much less cybersecurity information about ***specialized*** topics such as botnet infection and remediation.

Cybersecurity professionals know that user education is hard. If user education was easy and effective, many security risks would simply disappear: phishing for passwords and 4-1-9 advance fee fraud schemes, for example, would not work. Users would "get it" and not fall for those hoary old scams. But reportedly, users continue to cough up their credentials and continue to front tens of thousands of dollars to process the transfer of their "millions" from bogus Nigerian princes.

(b) We also need to recognize that botnets are inherently surreptitious and deceptive – they do their best to avoid being noticed by users. Given all the tricks that bots can pull, it can hard – even knowing that your system may be infected – to actually detect the presence of a bot, to say nothing of the problems that may arise when attempting to clean an infected system. (Remember, most security experts recommend nuking-and-paving rather than trying to remove malware in-situ).

(c) Many ISPs will insist their users get cleaned up. Enforcement of this policy can make it hard to tease apart effects associated with education and effects associated with policy enforcement by the ISP.

(d) Infection rates are already incredibly tiny – as previously mentioned in question 20, only 0.005% of all U.S. IPs are infected. Given that small base rate, it may be difficult to accurately measure the penetration and benefit of educational efforts.

VII. Incentives To Promote Voluntary Action To Notify Consumers

(22) Should companies have liability protections for notifying consumers that their devices have been infected by botnets? If so, why and what protections would be most effective in incentivizing notification? If not, why not? Are there other liability issues that should be examined?

As we have previously discussed in question 19, notification is not normally associated with liability. Liability issues are more likely to arise, if they do at all, when a provider attempts to help a user disinfect. Working on a malware infected system can sometimes result in the system crashing or becoming unstable, and it is not unusual for users to blame the computer technician who has been trying to assist them rather than the malware that has infected their system if that happens. By comparison, notification is a relatively low impact and low risk activity.

Risks in these and other situations are normally managed with disclaimers: "If you'd like help with your system, be sure you've completely backed it up before you bring it in for work, and please sign this liability waiver before we begin to work on your equipment."

If you would like to proactively incentivize notification, MAAWG would recommend ensuring that notifications are not a financial burden on firms. That is, if you want notifications to be routine, help pay the costs associated with creating and delivering those notifications; do not make botnet notifications yet another unfunded federal mandate.

(23) What is the state-of-practice with respect to helping end-users clean up their devices after a botnet infection? Are the approaches effective, or do end-users quickly get re-infected?

As we have previously mentioned, most security professionals will recommend that an infected system be nuked-and-paved, that is, formatted and reinstalled from scratch or restored from a clean backup.

If users do not choose to reformat and reinstall or are unable to execute that strategy (they cannot afford the cost, cannot find their original media, do not have backups of their files, etc.), consumer-grade remediation normally involves using one or more antivirus tools until no more malware is found. Yet even then, frustratingly, malware may still be present, albeit unknown to the antivirus software on the system.

If the ISP has not licensed a commercial antivirus product, users may be encouraged to try free antivirus products for that process. This recommendation needs to be made carefully, since some malware authors have released fake antivirus products that can actually infect a clean system or claim to find malware where none exists. It is the opinion of MAAWG that production, distribution or sale of fake antivirus malware should be a specifically cognizable federal offense carrying serious penalties, since it attempts to exploit weak users at their most vulnerable, and as such is a particularly heinous and despicable criminal act.

Users who are not comfortable attempting self-remediation may be referred to a commercial remediation provider partner, typically a service the user would pay for themselves. As we have previously discussed, the practical upper bound on the cost of such third-party work is effectively the cost of a replacement system plus the cost to migrate clean copies of critical files, less the resale value of the old infected system.

You also asked about users becoming re-infected: some do, some do not. Users who take the time to patch their systems and take other appropriate steps to harden their computers, as we have previously described, tend not to get re-infected. Users who do not take the time to harden their systems often do get re-infected.

(24) What agreements with end-users may need modification to support a voluntary code of conduct?

Most ISPs either have a formal contract for service with their customers or may publish terms of service (TOS) on their website. Depending on the content of a voluntary code of conduct, some terms of those contracts or TOS may require amendment, but it is impossible to say exactly without more information about what might be included in the voluntary code of conduct that would be ultimately forthcoming.

(25) Of the consumer resource scenarios described above, which would be most effective at providing incentives for entities to participate? Are there other reasons to consider one of these approaches over the others?

Most ISPs are economically rational entities. If you want them to participate in a voluntary code of conduct, ensure the program nets more revenue than it costs for participants. Help make a voluntary code of conduct a differentiator, or selling point, that might encourage a potentially reluctant ISP to actually sign up. For example, you might offer a Web badge or other consumer-visible incentive. Stress the fact that clean users are easier for the ISP to support, require fewer system and network resources, and are less likely to defect or be a source of "churn."

But recognize, ultimately, that while security used to be an "order qualifying" characteristic but not an "order winning" proposition, this is changing. That is, people may balk if security is not present, but interestingly more customers are putting the question, "How secure is the ISP?" much closer to the top of their shopping list. As leading U.S. ISPs move to an ever-improving level of baseline performance, security is becoming critical along with cost, performance and features.

(26) If a private sector approach were taken, would a new entity be necessary to run this project? Who should take leadership roles? Are the positive incentives involved (cost savings, revenue opportunity, etc.) great enough to persuade organizations to opt into this model?

Most American ISPs already have adopted solutions for managing botted customers, which is reflected in the low rate of infections in the United States. It is not clear that a new private sector organization is needed or that it would be successful if created.

(27) If a public/private partnership approach were taken, what would be an appropriate governance model? What stakeholders should be active participants in such a voluntary program? What government agencies should participate? How could government agencies best contribute resources in such a partnership?

Many public/private sector partnerships have been tried over the years. In the security space, two of the most widely known entities are probably the FBI's Infragard program and the various DHS critical infrastructure sector Information Sharing and Analysis Centers, or ISACs.

What we know from those experiments is that some Infragard chapters thrive, while others wither or become vestigial. Likewise, some ISACs are well subscribed and do important work, while others exist in name only.

The telling factors that seem to distinguish the programs that succeed from the programs that fail usually are:

- The partnership must have real/important work to do, work that is valued by the participants.
- A critical mass of participants must be available in the geographic area (an agriculture ISAC in Manhattan might be expected to fail, for example, given the limited number of urban farmers).
- Participants need to be trusted and prove worthy of that trust.
- Usable information-sharing channels (mailing lists, meetings, etc.) must exist.
- Information sharing needs to drive action.
- Information sharing needs to be bidirectional.
- Agency support must be genuine and not pro-forma.
- They need steady management with a light touch.

Governance, whether for a thriving or failing organization, commonly involves a governing board plus an executive director or co-executive directors. (In many respects, if the organization is otherwise sound, governance participants may find themselves with little work that needs to be done.)

You asked, "What stakeholders should be active participants in such a voluntary program?"

As a voluntary program, participation should be open to all trustworthy ISPs or individuals. There are nationality restrictions on Infragard membership that exclude many worthy individuals. Whether those entities elect to actively participate or not will be up to them.

You also asked, "What government agencies should participate? How could government agencies best contribute resources in such a partnership?"

It is not clear that any existing government agency is completely "right" for the envisioned activity.

- Remediation of botted end-user systems is not a core law enforcement mission, although law enforcement has a role in investigating and prosecuting bot herders and bot users.
- Participation by the intelligence community may make some non-governmental participants anxious or reluctant to share, and in most cases sharing of data from the intelligence community to unvetted participants would be difficult or impossible.
- US CERT might be a suitable participant, but they may not have sufficient resources, particularly if a distributed model is employed, to be effective without compromising higher priority responsibilities.
- The FTC mission is not practically aligned with tackling bots and their experience with tackling spam has shown they are under-resourced to tackle their existing portfolio.
- Historically, the FCC has regulatory responsibilities for some potential participating actors, but that regulatory history may negatively color, or provide an appearance of coercion, for what is meant to be a voluntary program.
- Perhaps the Office of Justice Programs might work? They have grant programs that are intended to leverage technology to fight crime (see http://www.ojp.gov/programs/technology.htm), although admittedly their normal focus is on justice-related entities, not ISPs.

When it comes to contributions from any agency that does participate, natural expectations will likely be:

- Resources (funding, facilities including meeting rooms and servers, staff flywheel support),
- Information (e.g., unclassified intelligence sharing), and
- Action on information shared from the private sector ("I've identified a botnet C&C and its located in Chicago...").

(28) If a government-run approach were taken, what government agencies should play leading roles?

See the discussion of agencies in the preceding question.

(29) Are there other approaches aside from the three scenarios suggested above that could be used to create a consumer resource and to incentivize detection, notification, and mitigation of botnets?

We would urge you to be clear on one potential customer population that might need help: there are some users who are botted, may have exhausted all free resources and are not able to afford commercial remediation assistance, and also not be able to afford to replace their infected system.

Where are those economically disadvantaged users to turn?

If we simply throw up our hands, and say, "Sorry, your infected system is your problem, not ours," we may come to regret that decision. If left infected, those botted systems can be turned against any or all of us: attacking government agencies, launching DDoS attacks against critical infrastructure, or being used as a stepping stone to hack defense industrial sites or even foreign countries.

Those economically disadvantaged users need help to get cleaned up. As in most things where all other resources have failed to successfully tackle a problem, the government may be all that is left, acting as a party of last resort, just as it does when it confronts the results of floods, hurricanes or earthquakes. This role, a "safety-net" function for the systems of our most economically-challenged citizens, would require distributed service centers, since most have little or no ability to transport themselves or their systems for help.

We also need to recognize that while we want and need to help our own citizens who have run out of options for cleaning up their infected systems, many of the botted hosts that routinely attack Americans are located ***outside our borders***. We also need a plan to help clean and correct those systems, not just the infected domestic ones.

A MAAWG Senior Technical Advisor has previously shared some thoughts on these requirements; see "We Need a Cyber CDC or Cyber World Health Organization" at http://pages.uoregon.edu/joe/ecrime-summit/ecrime-summit.pdf

(30) Are there other positive incentives that do not involve creation of an organized consumer resource that could encourage voluntary market-based action in detection, notification, and mitigation of botnets?

Consider offering a cash bounty program that would pay a reward for each new botted host that is identified and accurately reported, with evidence supporting the report. A bounty program of this sort would directly encourage private parties to identify and report the bots they encounter.

Also consider creating an application that users could run, which would voluntarily associate a contact address supplied by the user with the user's IP address. A blind remailer could then forward mail from authorized bot reporters to the email point of contact if a particular IP address appeared to be botted or otherwise have issues. This application would allow ISPs to be removed from acting as a critical notification conduit since it would no longer be necessary to connect an IP address having issues to a customer. The customer could be contacted directly, in a secure and privacy-preserving process.

Conclusion

Thank you for this opportunity for MAAWG to comment on " Models to Advance Voluntary Corporate Notification to Consumers Regarding the Illicit Use of Computer Equipment by Botnets and Related Malware." If you would like us to discuss any of our remarks in more depth, or if you have any questions, please do not hesitate to contact us.

Sincerely,

/signed/

Jerry Upton, Executive Director
Michael O'Reirdan, Chairman of the Board
Messaging Anti-Abuse Working Group (MAAWG)
jerry.upton@maawg.org
http://www.maawg.org

PPC2011-013