



March 2013

Final Report

U.S. Anti-Bot Code of Conduct (ABC)
for Internet Services Providers (ISPs)

Barrier and Metric Considerations

WORKING GROUP 7 - Botnet Remediation

Table of Contents

1	Results in Brief.....	3
1.1	Executive Summary	3
2	Introduction	4
2.1	CSRIC III Structure	5
2.2	CSRIC III WG7 Structure.....	5
2.3	Working Group 7 Team Members	5
2.4	Working Group 7 Deliverables:.....	7
3	Objectives, Scope, Methodology, and Recommendations.....	7
3.1	Barriers to ISP Participation	7
3.1.1	Objective	8
3.1.2	Scope.....	8
3.1.3	Methodology	8
3.1.4	Barriers to Code Implementation: Considerations.....	9
3.1.4.1	Technology Barriers.....	10
3.1.4.2	Consumer/Market Barriers.....	10
3.1.4.3	Operational Barriers	10
3.1.4.4	Financial Barriers	11
3.1.4.5	Legal/Policy Barriers.....	11
3.2	Metrics	12
3.2.1	Objective	13
3.2.2	Scope.....	13
3.2.3	Methodology	13
4	Conclusions & Recommendations	13
5	Acknowledgements	16
	Appendix 1 – Working Group 7 Mission Statement	17
	Appendix 2 – U.S. Anti-Bot Code of Conduct (ABCs) for Internet Service Providers (ISPs) Addressing Bot Activity in Broadband Networks	18
	Appendix 3 – Barrier Guide.....	35
	Appendix 4 – Metrics Guide.....	62
	Appendix 5 – Metrics Glossary	92
	Appendix 6 - Related Industry Security and Metrics Activity	96

1 Results in Brief

1.1 Executive Summary

The CSRIC III initially tasked Working Group 7, Botnet Remediation, with proposing a set of voluntary practices that would constitute the framework for a voluntary program for ISPs to follow to mitigate the botnet threat. In response, the U.S. Anti-Bot Code of Conduct for ISPs was delivered in March 2012 which addresses the threat of bots and botnets in residential broadband networks through voluntary participation.

As emphasized in the Code, “the growth of bot infected end-user devices represents a meaningful threat to the vitality and resiliency of the Internet and to the online economy.” Network security and customer access to the services provided by networks is of critical importance. Code participation offers a voluntary and collaborative approach that enhances participants’ ability to communicate with the public and their customers.

It is important to restate from the Code that constituents of the entire Internet ecosystem have important roles to play in addressing the botnet threat, and that ISPs depend on support from the other parts in the ecosystem. To maximize the effectiveness of anti-botnet efforts, it is essential to recognize the shared responsibilities that exist across the broad internet ecosystem. Such efforts, would involve active participation from such stakeholders as anti-virus and security vendors, applications and operating systems developers, device manufacturers, domain registrars and registries, end users, cloud-service providers, IT departments, search engines, web-owners, hosting services, and others. The FCC should support efforts to locate future botnet related activities in multi-stakeholder venues.

This report addresses aspects of the future work identified in the March 2012 U.S. Anti-Bot Code of Conduct for ISPs Final Report. First, it identifies potential barriers to participation in the Code, and provides guidance to evaluate or overcome specific barriers associated with each of the Code’s recommendations. Second, it addresses the current state and challenges associated with the development of standardized metrics for use in implementation of the Code. A guide is provided for both barriers to participation in the Code as well as for the ongoing development of metrics.

Recommendations for each report section, barriers to participation and metrics, are included in the Conclusion and Recommendations Section. Recommendations are made to treat the Code, Barriers Guide, and Metrics Guide as living documents which will require updates based on collaboration by Code participants with support from the Internet ecosystem community to address bots. WG7 further recommends the creation of a FCC sponsored botnet workshop involving members of the Internet ecosystem on addressing these recommendations.

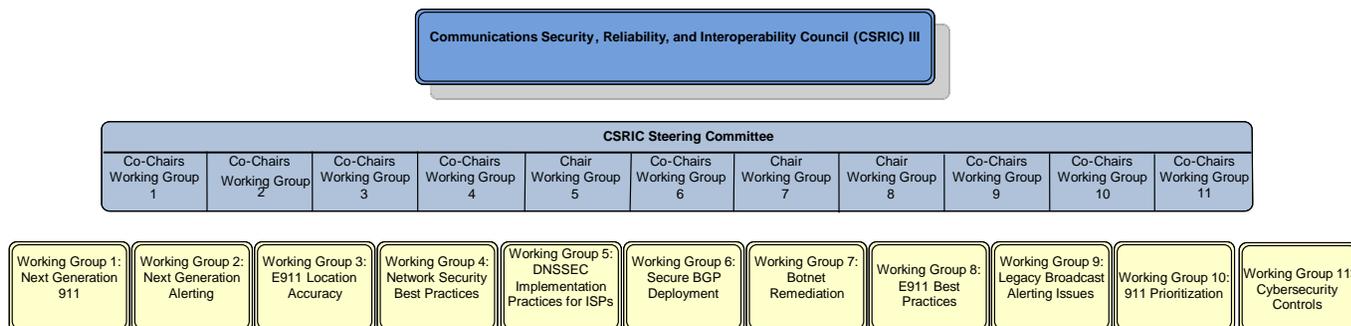
2 Introduction

CSRIC III established Working Group 7 (WG7) to address botnet remediation in broadband residential networks. In addition to creating the Code, WG7 was chartered to identify potential obstacles to ISP implementation of the U.S. Anti-Bot Code of Conduct for ISPs and steps the FCC could take that may help overcome these obstacles. WG7 was also requested to identify performance metrics to evaluate the effectiveness of the U.S. Anti-Bot Code of Conduct for ISPs at curbing the spread of bot infections.

Work to address these items and additional steps ISPs can take are documented in this report. A specific guide is provided for each of the sections: A Barriers Guide addressing barriers to Code participation as well as a Metrics Guide for the ongoing development of metrics. The guides will provide useful tools for ISP participation in the Code and provide templates for contribution of future work. ISPs of all sizes should find tangible and practical steps to help in their efforts to combat bots.

The report concludes with recommendations related to participation in anti-botnet initiatives and efforts to overcome barriers to Code participation.

2.1 CSRIC III Structure



2.2 CSRIC III WG7 Structure

WG7 is chaired by Michael O'Reirdan, Co-Chairman of the Messaging Anti-Abuse Working Group (M³AAWG¹), and vice-chaired by Dr. Peter Fonash, Chief Technology Officer, Office of Cybersecurity and Communications, Department of Homeland Security. Robert Mayer, Vice President – Industry and State Affairs, USTelecom Association, facilitated the development of the Barriers to the Code Implementation Guide. Members of WG7 include representatives from ISPs, suppliers of software and network equipment, academia, as well as other organizations that are a part of the Internet ecosystem.

2.3 Working Group 7 Team Members

Working Group 7 consists of the members listed below.

Name	Company
Michael O'Reirdan – Chair	M ³ AAWG
Peter Fonash – Vice-Chair	Department of Homeland Security
Neil Schwartzman - Secretary	CAUCE
Robert Thornberry – Editor	Bell Labs, Alcatel-Lucent
Paul Diamond – Editor	CenturyLink
Vernon Mosley – Liaison	FCC
Michael Little	Applied Communication Sciences
Alex Bobotek	AT&T
John Denning	Bank of America

Uma Chandrashekar	Bell Labs, Alcatel-Lucent
Chris Lewis	CAUCE, Spamhaus
Michael Glenn	CenturyLink
Jay Opperman	Comcast
Matt Carothers	Cox
Gunter Ollmann	Damballa
Brian Done	Department of Homeland Security
Daniel Bright	EMC INC
Mats Nilsson	Ericsson
Kurian Jacob	FCC
Bill McInnis	IID
Chris Sills	IID
Tim Rohrbaugh	Intersections
Barry Greene	ISC
Merike Kaeo	ISC
Ed White	McAfee
Kevin Sullivan	Microsoft
Matthew Tooley	National Cable & Telecommunications Association
Jon Boyens	National Institute of Standards and Technology
Craig Spiegle	Online Trust Alliance (OTA)
Bill Smith	PayPal
Gabe Iovino	REN-ISAC
Johannes Ullrich	SANS Institute
Adam O'Donnell	Sourcefire
Alfred Huger	Sourcefire
Kevin Frank	Sprint
Greg Holzapfel	Sprint
James Holgerson	Sprint
Michael Fiumano	Sprint
Maxim Weinstein	StopBadware
Patrick Gardner	Symantec
John Griffin	Telecommunication Systems Inc.
Chris Roosenraad	Time Warner Cable
Tice Morgan	T-Mobile
Joe St Sauver	University of Oregon/Internet2
Robert Mayer	USTelecom
Eric Osterweil	Verisign
John St. Clair	Verizon
Timothy Vogel	Verizon

Table 1 - List of Working Group 7 Members

2.4 Working Group 7 Deliverables:

1. U.S. Anti-Bot Code of Conduct for Internet Service Providers: March 22, 2012
2. Barriers to Code Participation and Metrics: March 6, 2013

3 Objectives, Scope, Methodology, and Recommendations

The following sections address Working Group 7 efforts regarding Barriers to Code Participation and Metrics. The Problem Statement being addressed is identified in the ABCs for ISPs Code (attached) as follows:

Problem Statement: The growth of bot infected end-user devices represents a meaningful threat to the vitality and resiliency of the Internet and to the online economy. Botnets are networks of Internet-connected end-user computing devices infected with bot malware, which are remotely controlled by third parties for nefarious purposes. Bots and botnets can lead to theft of personal information, attacks against public and private networks, and exploitation of end-users' computing power and Internet access.

In order to reduce the bot infections in residential end-user devices and mitigate the potential exploitation of the bots, the voluntary U.S. Anti-Bot Code of Conduct for ISPs was developed by the Working Group 7 members.

3.1 Barriers to ISP Participation

Since the inception of Working Group 7, including subsequent efforts to produce the initial March 2012 report and this latest report, it has been clear to the participants that many of the recommendations that were adopted involve varying levels of commitment and complexity. It was also understood that private sector participants compete internally for investment capital and human resources and that decisions to prioritize such investments are increasingly dependent upon providing a solid business case to the decision-makers. This is particularly challenging in the current economic environment when carriers and ISPs face a host of threats in both the physical and virtual realms. Furthermore, Code activities can only happen with participation from other Internet ecosystem participants and end users, such as creation of software patches to address exploits and malware, and end-user system remediation.

The Barriers Guide is designed to encourage broader participation by ISPs by providing a structured set of resources related to specific recommendations including the best available guidance on implementation. It is our expectation that the guidance will evolve over time as more organizations assess the impact of the problem and the benefits associated with implementing specific actions.

As participants in the broader Internet ecosystem, ISPs have a significant interest in working with other key participants and adopting policies to address the botnet challenge. Numerous

barriers, many of which can be removed, currently limit the steps ISPs and other stakeholders can take to evaluate and remedy threats affecting their networks. Such steps include the adoption of valuable measures and best practices like those recommended by the voluntary U.S. Anti-Botnet Code of Conduct for ISPs. The barriers to broad adoption of the Code range from evident concerns about costs, funding, and marketplace dynamics, to complex policy considerations concerning access rights, information sharing, as well as the scope and methods of cooperation between ecosystem stakeholders, governmental entities, end users, and other potential parties.

Identifying and resolving these concerns is a critical next step toward the widespread adoption of best practices by ISPs.

3.1.1 Objective

In order to reduce the spread of bot infections and facilitate the adoption of the Code of Conduct, WG7 was chartered to identify potential barriers to implementation of the U.S. Anti-Bot Code of Conduct for ISPs and steps the FCC can take that may help overcome these barriers.

As a result, a Guide to Addressing Barriers to Implementation of the Code was created (see Appendix 2) as an aid to ISPs to help identify potential barriers to Code participation and provide information and references to address them.

3.1.2 Scope

The CSRIC Working Group 7 created the Barriers to Code Participation sub-group as part of future work identified in the March 2012 Final Report. The group was tasked with examining barriers that may exist to ISP Code participation and to provide guidance on how to address them based on information and evidence available at this point in time.

In order to facilitate broader ISP participation, barriers to participation were identified and guidance is provided to evaluate or overcome specific barriers associated with each of the 10 recommendations that fall within the five principal areas of activity (Education, Detection, Notification, Remediation, and Collaboration).

3.1.3 Methodology

ISPs have considerable flexibility in implementing various activities recommended by the Code. Working Group 7 began by considering the following points:

i) *Why Discuss Barriers to Adoption?*

Answer: By articulating barriers to adoption, it is hoped that future processes can be put in place, and resources developed, to overcome those barriers, and thus increase the likelihood of more widespread Code participation by US ISPs.

ii) *You Can't Manage What You Can't Measure.*

Answer: One barrier may relate directly to the community’s problems when it comes to measuring bots. Without consistent objective agreed upon industry based measurements, ISPs may find it difficult or impossible to tell if they have a bot problem, and if so, whether efforts to correct it will have, or have had, any material effect. Until we can overcome this limitation, exhortations to participate in the Anti-Bot Code will have to be reliant on reasons other than metrics and other empirical evidence. Metrics are addressed in Section 3.2.

iii) *Fighting Bots Requires Broad Participation; It Can’t Be Just “An ISP Thing.”*

Answer: Botnets are and fuel criminal enterprises. In order for work against bots to be effective, successfully curtailing bots or mitigating their impact will require collective action by all parts of that ecosystem, from such stakeholders as anti-virus and security vendors, applications and operating systems developers, device manufacturers, domain registrars and registries, end users, cloud-service providers, IT departments, search engines, web-owners, hosting services, and others. If anti-bot efforts lack that broad base of participation, criminals will be able to continue to compromise systems and use those botted hosts for nefarious purposes. In particular, note that ISPs cannot directly manage consumer devices, which means that end-user cooperation is critical when it comes to remediating botted machines.

iv) *That Said, ISPs Do Have A Critical Role to Play In Fighting Bots.*

Answer: The Code sought to lay the groundwork for future coordination among various stakeholders by defining a set of actions appropriately directed to the limited but key role ISPs play in addressing this important issue. This discussion of barriers furthers that goal, by more specifically articulating issues that ISPs may face and need to overcome as part of their participation activities, and thus more clearly illustrate for other ecosystem participants areas in which their unique expertise and capabilities, are essential.

v) *The Cost Of Fighting Bots Can Be One Barrier to Participation, But It Isn’t The Only One.*

Answer: The barriers identified in this report are not just related to costs. Of course cost is an issue in implementing many activities, and a reasoned cost-benefit analysis is the touch-stone of any security risk assessment and risk mitigation. There are likely to be different ways to overcome many of the barriers identified, and it is hoped that by enumerating these barriers, the creative resources of the entire Internet ecosystem can be brought to bear on finding cost-minimizing and cost-effective solutions for overcoming them.

3.1.4 Barriers to Code Implementation: Considerations

Code participation can pose barriers to the extent that the requisite activities impact existing methods and procedures and the resources of multiple organizations within an ISP. ISPs will need to understand not only what changes may be required, but which organizations would be impacted from a process, resource, and budget perspective. Consideration must be given to scalability and the levels of inter-departmental integration as well as any on-going support required to implement a recommendation. In addressing barriers, efforts were made to categorize each identified barrier into one of the following categories:

3.1.4.1 Technology Barriers

Technology barriers are barriers where current technical solutions may be insufficient to tackle the botnet threat or where those solutions may have other unacceptable side effects. The significance of technology solutions as barriers depends on the degree to which technology is required to implement a specific recommendation. Some solutions may require minimum technical resources (assets and people) while others might be more demanding including levels of systems integration. Implementing a recommendation is also likely to depend on the current “as-is” situation within an ISP including internal capabilities, resources and priorities or access to third-party resources. Technology barriers may directly translate into financial barriers. A technical solution may exist but the cost to implement it in the ISP may be prohibitive.

3.1.4.2 Consumer/Market Barriers

Customer or Market Barriers are barriers that arise out of implementing solutions that may be viewed by customers as ineffective (e.g., customers choosing not to participate) or undesirable (e.g., privacy, restrictive terms and conditions). ISP activities may also have consequences in the marketplace by raising costs for their products or concerns about a provider’s vulnerabilities in relation to other competitors. When considering barriers related to consumers or market impacts, it is important to understand how a particular change will impact individual customers. Implementations that have a negative customer experience can directly impact ISP revenue. To the extent that some ISPs are marketing their capabilities in the area of bot protections, it would be useful to understand how market dynamics are currently impacted.

3.1.4.3 Operational Barriers

Operational barriers are barriers that could have an unacceptable impact on an organization’s primary mission and resources. Organizations that have operational responsibilities are typically held accountable through well-defined performance measurements. (e.g., average transaction time for customer care representatives). Any new operational practice must be implemented with sufficient attention to:

- development and updated operational methods and procedures;
- reallocation of existing or additional resources;
- scalability of solutions;
- operational performance goals, critical success factors and ability to measure results.

3.1.4.4 Financial Barriers

Financial barriers result from any inability to quantify costs or benefits associated with implementing specific recommendations. In the current economic environment, the absence of a company-specific business case may be a significant barrier to adoption. Private sector companies have a fiduciary responsibility to make decisions based on the prudent allocation of capital. This does not mean that every investment has to have a demonstrated return on investment. There are many instances where companies contribute resources in ways that are sometimes captured as “good will” or to serve some larger public or private interest. Still, when it comes to deployment of investments in security-related technologies, most companies rely on quantifiable data to prioritize and allocate new investment in systems, processes, and human resources. In the case of the U.S Anti-Botnet Code of Conduct, some recommendations require more one-time or on-going capital and operating expenses. Generally speaking, the larger the investment, the greater the expectation by management for rigorous cost-benefit analyses.

3.1.4.5 Legal/Policy Barriers

Laws that discourage collaboration and information-sharing among ISPs, either in the U.S or abroad, can give cover to malicious actors.² Malicious actors have the ability to disguise botnets as legitimate traffic, which places security professionals in the untenable position of potentially intruding upon an end-user’s expectation of privacy or allowing dangerous and costly criminal activity to take place.³ Furthermore, ISPs are not always able to stop traffic even if it is recognized as malicious. For example, the traffic could be interwoven with commercial activity, such that the cost of mitigation would be greater than the harm posed by the botnet. So long as they exist, botnets will increase the cost of doing business online.⁴

Guide to Implementation: ISPs have considerable flexibility in implementing various activities recommended by the Code. In an effort to help ISPs evaluate specific recommendations and the barriers and additional considerations associated with implementation, the Botnet Working Group developed a guide to help with assessment and implementation (See Appendix 3). The Guide lists each of the five areas in the Code and the accompanying 10 recommendations. The Code was designed to encourage participation in at least one activity in each of the five areas. The 10 recommendations found in the Guide vary with respect to implementation requirements, costs, and potential benefits. To the extent available, the Guide contains information based on successful implementation for each of the recommendations. As a result, individual ISPs, based on their own unique circumstances can make informed decisions as to the value of each recommendation. To promote and accelerate adoption, the Guide identifies and highlights a set of options that minimize the burden on ISPs. It is our expectation that as new information and evidence is presented regarding the reduction or elimination of specific barriers, ISPs will find it

² See *id.*

³ See generally, Can We Beat Legitimate Cyber Behavior Mimicking Attacks from Botnets?, IEEE (2012).

⁴ See, e.g., Department of Commerce, White House Announces Public-Private Partnership Initiatives to Combat Botnets, <http://www.commerce.gov/news/press-releases/2012/05/30/white-house-announces-public-private-partnership-initiatives-combat-b> (May 30, 2012).

in their interest to implement more advanced solutions that are likely to be more effective in combating the botnet threat.

3.2 Metrics

"If you cannot measure it, you cannot improve it." – Sir William Thomson Kelvin (Lord Kelvin) Scientist, (1824- 1907)

In addition to the barriers discussed in Section 3.1 another key barrier to effective Code participation is the current inability to uniformly measure the bot population and the results of activities to reduce bots. Without consistent objective agreed upon industry based measurements, ISPs may find it difficult or impossible to tell the extent of the bot problem on their network, and if so, whether efforts to correct it will have, or have had, any material effect. Until we can overcome this limitation, exhortations to participate in the Anti-Bot Code will have to be reliant on reasons other than metrics and other empirical evidence.

As the bot problem is not a closed system, we caution that outcome based metrics alone cannot be used to adequately evaluate activities of the Code. In other words, ISPs could fully participate in Code activities but not necessarily see any significant reduction in bots on their networks. Instead, we recommend a combined methodology of outcome and activity based measures. For example, a measurement system could seek to measure both the total number of detected infected subscribers and the number of infected subscribers notified by the ISP. These are outcome and activity based metrics, respectively.

In some cases, Internet ecosystem participants other than the ISP may notify the end user of a bot infection. A bank, for example, may detect an end-user's computer infected with malware and provide immediate notification to the customer of the infected device. The customer may be more inclined to address the problem immediately in this scenario than through a notification by their ISP. Information sharing between the bank and ISP may be important to understand if the customer addressed the problem and to ensure duplicate notifications are not sent, possibly creating a negative user experience.

These combined metrics can be used to gauge the overall efficacy of voluntary Code participation with the proper balance of outcome and activity based measures. In this group's initial report we identified the need for the Code to be able to adapt to changing online threats. An effective metrics approach can help guide that adaptation in future iterations of this Code. Metrics can also benefit the ISPs who participate in the Code. With a common metrics system they can make objective comparisons of different approaches, compare themselves to aggregate benchmarks, and help determine the cost benefit analysis of solutions.

The working group agrees that effective implementation of the Code will require a common and agreed to metrics vernacular. Common terms and definitions for metrics are essential for communication of goals and progress toward implementation of the specifics of the Code. This section describes the working group's investigation of such a system, analysis of existing efforts, and analysis for a comprehensive solution.

3.2.1 Objective

The CSRIC tasked Working Group 7, Botnet Remediation, with identifying performance metrics to evaluate the effectiveness of the U.S. Anti-Bot Code of Conduct for ISPs at curbing the spread of botnet infections.

Metrics need to be defined to measure the impact of ISP Code participation in reducing the number of bot infections.

3.2.2 Scope

While our focus on scope is on the ISPs participating and implementing the Code we recognize that in order to measure the full impact of the Code we will require a broader participation by other ecosystem participants.

We found that there are two types of metrics required:

- 1) Metrics to measure overall effectiveness of Code participation.
E.g., Reduction of bot infections over time.
- 2) Activity based metrics
E.g., How many notifications were delivered in the reporting period and what was their effectiveness (combining this metric with #1).

Additionally, metrics could be part of a qualitative case study and enable consistent bilateral comparison of effectiveness between participants.

3.2.3 Methodology

Working Group 7 quickly recognized the need to take a broad approach to gathering and developing metrics and thus set the following three point plan:

- 1) Create a Bot Metrics Guide to analyze bot metric issues and make recommendations toward development of consistent, obtainable, and meaningful bot metrics. Refer to Appendix 3.
- 2) Create a glossary of bot terms to ensure consistent terminology in analyzing and investigating bot metrics. Refer to Appendix 4.
- 3) Examine current industry bot metrics activity to identify synergies with the WG7 metrics effort. Refer to Appendix 5.

4 Conclusions & Recommendations

WG7 was charged by CSRIC members with proposing a set of agreed-upon voluntary practices that would constitute a framework for a voluntary implementation model for ISPs to follow to mitigate the botnet threat. In March 2012, WG7 delivered the U.S. Anti-Bot code of Conduct for ISPs to address the threat of bots and botnets in residential broadband networks. Furthermore, the working group was charged with identifying potential ISP implementation barriers to the

Code, and identifying steps the FCC can take to help overcome these barriers. The working group was also charged with identifying performance metrics to evaluate the effectiveness of the Code at curbing the spread of bot infections. This March 2013 Final Report includes the Code from the March 2012 report as well as addresses the barriers to ISP implementation of the Code and Code effectiveness metrics.

WG7 advocates that the Code, Barriers Guide, and Metrics Guide be living documents which will require updates based on experiences shared through the collaboration of Code participants. In order for the bot problem to be adequately addressed, WG7 believes that the bot work of the group will need to be expanded to the rest of the Internet ecosystem so that synergies with the Code can be established and issues worked with a broader view of the problem (e.g., website owners who might detect anomalous activity from their clients and antivirus vendors who have end node intelligence on bots) and solution space (e.g., Internet ecosystem multi-stakeholder solutions beyond actions by ISPs). Lastly, WG7 believes in the urgency to act in defending residential broadband networks from bot infections, the necessity to address the bot problem with an Internet ecosystem multi-stakeholder approach, and the need for a continued focus on bot reduction and mitigation in order to leverage the effectiveness of actions taken by the Internet ecosystem community to reduce the spread of bot infections.

In recognition of the above, WG7 offers the following recommendations:

Barriers to Code Participation:

1. WG7 recommends the FCC working in partnership with other federal government agencies and industry facilitate ISP awareness of the Code Barriers Guide and encourage ISPs to use the guide as a resource in evaluating and planning Code participation. FCC promotion can take the form of press releases, FCC website coverage, sponsored workshops, and outreach briefings. Further, WG7 recommends that the FCC support the Barriers Guide as a living document and facilitate efforts by all members of the Internet ecosystem in creating updates by charging future CSRICs to evaluate the roles and contributions of all ecosystem participants in any future updates to the Code Barriers Guide.

Metrics:

2. WG7 recommends that the FCC working in partnership with other federal government agencies and industry facilitate the creation of case studies on bot mitigation activities, to examine metrics created around particular bot remediation efforts, by identifying up to three prolific bot-infection events occurring over the next year, and seeking participation in the case studies from the multi-stakeholder community. Case Studies focus on specific events and provide a detailed analysis which can be used to better understand what is important to measure. These efforts will need to involve not only the ISP community but the larger Internet ecosystem as well.

A study currently in progress at a major US university in the south east to be published in

March 2013 on ISP actions related to DNSChanger should be considered by the FCC as an initial case study. This effort, still in progress at the time of this writing, centers on the DNSChanger malware and related customer notification methods used by ISPs. This approach may well show the cooperative, collaborative steps required for the future development of overarching metrics involving multiple ISP approaches, and could be used as a model for future case studies. These steps are a microcosm of the recommendations in the Code.

3. We further recommend the FCC working in partnership with other federal government agencies and industry leverage industry sponsored pilot programs, e.g., M³AAWG Bot Metrics Pilot Program, to examine the collection and sharing of metrics around particular bot efforts, based on and incorporating results of the case studies, by identifying metrics to be collected and shared across the Internet ecosystem to help reduce bot infections. Pilot Programs differ from Case Studies in that they collect metrics on an on-going basis and provide participants with anonymized data which can baseline relative performance. Such programs are required to test which metrics may be useful and obtainable and which are not. Some metrics may, in fact, not be feasible to obtain at all. For comparative purposes, the metrics definitions must be reasonably standardized between ISPs. We anticipate that some metrics methods used by ISPs will lend themselves to comparative analysis and some will not. Participation in pilot programs will indicate which, if any, are viable.
4. We recommend that the FCC working in partnership with other federal government agencies and industry facilitate research in bot metric development by identifying gaps and shortfalls between outcome-based metrics needed to measure the effectiveness achieved by ISPs participating in the Code and metrics ISPs typically collect and share across the Internet ecosystem related to bot mitigation strategies. This research would seek to develop common definitions of bots for industry and methods to distinguish bots and bot families from other malware for measurement purposes. The research would provide input to industry efforts to define a set of implementable, outcome-based effectiveness metrics and standard methods of measurement based on best practices, case studies, and pilot programs.

Workshop:

5. We further recommend that the FCC working in partnership with other federal government agencies and industry establish a vehicle such as a workshop or other education technique (i.e., a webcast) to foster ongoing dialogue around these issues. We encourage the inclusion of appropriate international participants (e.g., representatives from bot remediation bodies from Australia, Japan, etc.), as practical, in order to incorporate their learnings. WG7 believes this approach, incorporating and discussing the results of the above recommendations, will lead to further recommendations on Internet ecosystem multi-stakeholders approaches to best contain the spread of bot infections.

WG7 believes by expeditiously taking voluntary action on the above recommendations, the FCC will significantly contribute to, and facilitate the development and implementation of voluntary practices that can be followed by the Internet ecosystem multi-stakeholders to combat the spread of bot infections.

5 Acknowledgements

The Chairman of Working Group 7 would like to acknowledge the following members of the Working Group for their particular efforts:

Paul Diamond of CenturyLink and Robert Thornberry of Bell Labs, Alcatel Lucent for their tireless efforts as editors for the work products of the Working Group.

Robert Mayer of USTelecom for leading the team that produced the Barriers to Code Participation component of the report. Tim Vogel of Verizon for his dedicated work on the Barriers Guide.

Kevin Sullivan of Microsoft for leading the Metrics team and Joe St Sauver of the University of Oregon/Internet2 for his work on producing both the Metrics Guide and Glossary.

I would like to thank Dr. Peter Fonash of the Department of Homeland Security and Vernon Mosley of the Federal Communications Commission for their considerable efforts in facilitating the production of the report.

Finally I would like to thank all the members of the Working Group for their time and effort.

Appendix 1 – Working Group 7 Mission Statement

Appendix 2 – U.S. Anti-Bot Code of Conduct (ABCs) for Internet Service Providers (ISPs)
Addressing Bot Activity in Broadband Networks

Appendix 3 – Barrier Guide

Appendix 4 – Metrics Guide

Appendix 5 – Metrics Glossary

Appendix 6 – Related Industry Security and Metrics Activity

Appendix 1 – Working Group 7 Mission Statement

Working Group 7 – Botnet Remediation

Chair – Michael O’Reirdan, Chairman, Messaging Anti-Abuse Working Group

Description: This Working Group will review the efforts undertaken within the international community, such as the Australian Internet Industry Code of Practice, and among domestic stakeholder groups, such as IETF and the Messaging Anti-Abuse Working Group, for applicability to U.S. ISPs. Building on the work of CSRIC II Working Group 8 ISP Network Protection Practices, the Botnet Remediation Working Group shall propose a set of agreed-upon voluntary practices that would constitute the framework for an opt-in implementation model for ISPs. The Working Group will propose a method for ISPs to express their intent to opt-into the framework proposed by the Working Group.

The Working Group will also identify potential ISP implementation obstacles to the newly drafted Botnet Remediation business practices and identify steps the FCC can take that may help overcome these obstacles.

Finally, the Working Group shall identify performance metrics to evaluate the effectiveness of the ISP Botnet Remediation Business Practices at curbing the spread of botnet infections.

Appendix 2 – U.S. Anti-Bot Code of Conduct (ABCs) for Internet Service Providers (ISPs) Addressing Bot Activity in Broadband Networks

(see <http://transition.fcc.gov/bureaus/pshs/advisory/csric3/CSRIC-III-WG7-Final-Report.pdf>)

7 Appendix

U.S. Anti-Bot Code of Conduct (ABCs) for Internet Service Providers (ISPs) Addressing Bot Activity in Broadband Networks Final March 22, 2012

1. Introduction

The growth of bot* infected end-user* devices represents a meaningful threat to the vitality and resiliency of the Internet and to the online economy. Note that bot infection and bot are used synonymously in this document to refer to an end-user device infected with bot malware. Botnets are networks of Internet-connected end-user computing devices infected with bot malware*, which are remotely controlled by third parties for nefarious purposes.

Bots and botnets can lead to theft of personal information, attacks against public and private networks, and exploitation of end-users' computing power and Internet access. Public awareness of bots, their impact, and the resulting security and privacy issues is low. This voluntary code of conduct ("Code") provides a set of principles and recommended activities that Internet Service Providers may adopt to help address the threats presented by the presence of bots and botnets in residential broadband networks.

It should be recognized that bots impact the entire Internet ecosystem* and that successfully curtailing bots or mitigating their impact will require collective action by all parts of that ecosystem, including end-users, software developers, search providers, websites, e-commerce sites, and others. End-user devices are outside of the control of ISPs* hence all participants in the Internet ecosystem need to work together to address this issue. This Code seeks to lay the groundwork for future coordination among various stakeholders by defining a set of actions appropriately directed to the limited role ISPs can play to help address this important issue. The Code recognizes the substantial variability in size, resources, business models and environments, expertise, and abilities of the ISPs in the United States. The success of ISP activities relies on similar efforts by other Internet stakeholders.

The core requirements for participation in this Code are set forth in Section 5. The other sections of this document contain background information or additional explanatory material.

* Definition found in Glossary

2. Definitions of Key Terms

Note to the Reader:

Any discussion of bots inevitably involves a unique technical vocabulary. Recognizing that many readers may not be familiar with some of those specialized terms, the Code includes a glossary as Appendix 2. Any term appearing in the glossary will be marked with an asterisk "*" in the body of the Code text the first time it appears as a way of alerting the reader that a definition is available in the glossary.

3. Objectives and Principles

a. The objectives of this Code are to:

1. Provide an initial framework for ISPs to better understand and help address the bot issue; and
2. Encourage ISPs to
 - Educate end-users of the threat posed by bots and of actions end-users can take to help prevent bot infections;
 - Detect bot activities or obtain information, including from credible third parties, on bot infections among their end-user base;
 - Notify end-users of suspected bot infections or help enable end-users to determine if they are potentially infected by bots; and
 - Provide information and resources, directly or by reference to other sources, to end-users to assist them in remediating bot infections.

b. Implementation of the Code will be guided by the following principles:

1. Voluntary — participation is voluntary and encourages types of actions to be taken by ISPs, however this Code does not require any particular activity.
2. Technology neutral — this Code does not prescribe any particular means or methods.
3. Approach neutrality — this Code does not prescribe any particular approach to implement any part of this Code.
4. Respect for privacy — ISPs must address privacy issues in an appropriate manner consistent with applicable laws.
5. Legal compliance — activities must comply with applicable law.
6. Shared responsibility — ISPs, acting alone, cannot fully address the threat posed by bots. Other Internet ecosystem participants must also do their part.
7. Sustainability — ISPs should seek activities that are cost-effective and sustainable within the context of their business models.

8. Information sharing — ISPs should indicate how they are participating in the Code and share lessons-learned from their activities with other appropriate stakeholders. All information sharing between ISPs and other involved parties must be performed in accordance with applicable laws including, but not limited to, antitrust and privacy laws.
9. Effectiveness — ISPs should be encouraged to engage in activities that have been demonstrated to be appropriate and effective.
10. Effective Communication — Communication with customers* should take into account various issues such as language and make sure that information is provided in a manner that is reasonably expected to be understood and accessible by the recipients.

4. Scope and Roles

This Code was drafted specifically for ISPs and other service providers offering broadband Internet access service to residential end-users. The activities in this Code may be adapted for use by other Internet providers and participants.

This Code is not meant to be an all-inclusive approach to online security, but is meant to coexist with other current and future efforts. It anticipates a significant role for other Internet ecosystem participants, including but not limited to:

- Security software vendors
- Operating system developers
- End-user focused organizations
- Providers of Internet content, applications, and services

Online security should include a multifaceted, flexible approach using advice and tools from various reputable sources.

a. Definition of Success

Initial success of this Code will be assessed in terms of participation by the ISP community. Support by the Internet ecosystem at large, however, is seen as paramount to the ultimate success in the fight against bots.

b. Benefits of Participation in the Code

The following high-level benefits may result from meaningful ISP participation in this Code:

- Increased security of end-user information and devices and for U.S. infrastructure;
- Increased awareness of the bot threat and how to address it among end-users, ISPs, and other Internet-related industry participants;

* Definition found in Glossary

The Communications Security, Reliability and Interoperability Council
Final Report

Working Group 7
March, 2012

- Notification* and remediation* of bot activity on bot infected end-user devices;
- Creation of an environment in U.S. residential broadband networks that is even more hostile to the deployment and utilization of bots; and
- Development and wider use of effective notification and remediation architectures and tools across end-users and ISPs.

Some ISPs participating in the Code development process who have previously implemented some aspects of the Code have experienced beneficial results in areas such as lower call volumes to help desks from customers with infected machines, reduced upstream bandwidth consumption by denial-of-service attacks and spam*, increased customer goodwill and lower customer churn, and reduction in spam-related complaints from other ISPs. Although individual results may vary, ISPs are encouraged to look for specific ways in which Code participation bolsters their overall broadband business, and to share those experiences with other ISPs. In addition, ISP participation in this Code may enable ISPs to generate tangible metrics relating to the impact of specific activities on the ISPs overall broadband business operations, which in turn may support development or deployment of further anti-bot activities.

* Definition found in Glossary

5. Parameters for Participation

Participation in this Code is voluntary.

Voluntary Code of Conduct Participation Requirements

To participate in this Code, an ISP will engage in at least one activity (i.e., take meaningful action) in each of the following general areas:

- Education - an activity intended to help increase end-user education and awareness of botnet issues and how to help prevent bot infections;
- Detection - an activity intended to identify botnet activity in the ISP's network, obtain information on botnet activity in the ISP's network, or enable end-users to self-determine potential bot infections on their end-user devices;
- Notification - an activity intended to notify customers of suspected bot infections or enable customers to determine if they may be infected by a bot;
- Remediation - an activity intended to provide information to end-users about how they can remediate bot infections, or to assist end-users in remediating bot infections.
- Collaboration - an activity to share with other ISPs feedback and experience learned from the participating ISP's Code activities.

The concept of engaging in "at least one activity" in each of these general areas is intended to encourage some level of activity in each of the five areas noted above as part of an overall nationwide process of creating an environment in U.S. residential broadband networks that is even more hostile to the deployment and utilization of bots. It is intended to support and encourage a wide range of flexible efforts to experiment and innovate with various methods of education, detection*, notification, and remediation. In that same vein, the requirement to share feedback with other ISPs is not intended to dictate any specific means or methods of sharing such feedback.

* Definition found in Glossary

6. End-User Education

a. Overview

End-users are ultimately responsible for protection of their devices and for remediating an infected device. ISPs, like many other Internet participants and government actors, can assist in helping to educate end-users about the threats presented by bots and the steps end-users can take to protect their devices and remediate infections.

b. Recommended Action:

1. Education about bot prevention*:

ISPs should make available information about the prevention of bot infections and related issues. At a minimum, such information should include:

- How and why end-users must keep their software updated for computers and devices with readily available software updates.
- The importance of using effective and current security software from a reputable vendor.
- The importance of backing up user data and software and how to do it effectively.
- Basic end-user actions for minimizing exposure to bot infections while using the Internet.

It is expected that many ISPs will be able to accomplish this goal by providing this information directly to their subscribers or by linking to existing, publicly-available sources of such information.

2. Support of end-user bot remediation efforts:

Along with information on prevention, ISPs should make available (e.g., via ISP publications, third-party publications or web links) information on how end-users can generally remediate bot infections. In this area, it is expected that ISPs will be able to accomplish this goal by linking to existing, publicly available sources of such information or by creating new sources of such information, either individually or in conjunction with others.

In connection with an ISP's end-user notification activities, ISPs should include in such notices or otherwise, information on where the recipient might turn for additional information and assistance. Such information might include links to publicly-available online information, security tools, or suggestions to seek assistance of a computer professional. Additional subjects and references that an ISP might wish to include are:

* Definition found in Glossary

- Risks to the end-user and Internet community from using a device that is believed to be infected,
- How to identify and remove common forms of bot infections,
- Publicly-available tools or services (free or paid) to assist in the detection and removal of bot infections, and
- Guidance on where to find additional (free or paid) assistance.

3. Guidelines:

In addressing the above requirements, ISPs should consider these guidelines:

- Offer educational information and resources directly or through referral to third party services.
- Keep educational content concise and focused on the most important things users need to know.
- Ensure that instructions can be followed by an audience of non-technical users.
- Use multiple media, e.g., images, videos, text, captions, etc., and, where helpful, multiple languages to maximize customer understanding and accessibility.
- Help end-users determine if they have a bot infection by providing information or pointing to resources that describe anomalous behaviors of bot infected devices and the availability and use of bot detection software tools or services.

7. Bot Detection

a. Overview

As bots evolve, so must the tools and techniques used for detecting them. The challenge to detection lies in the versatility that bot traffic has achieved to avoid many singular techniques used for detection mechanisms, such as simple pattern matching. Detection may be complicated by the fact that some Internet applications, like distributed host-based caching content delivery networks, online gaming applications, and other such services may exhibit behavior similar to that of malicious bots, and may utilize similar technologies. ISPs should employ care in identifying impacted parties for notification and remediation.

b. Recommended Action:

ISPs can find out about malicious activity and bot compromised end-user devices in a variety of ways:

1. Receiving notifications from external entities, particularly those designed to aid with the overall understanding and real-time dissemination of bot related data. A list of resources is listed in Appendix 2.
2. Deploying capabilities within their networks that aid in identifying potential bot infections.
3. Directing customers to tools, a web portal, or other resources that enable customers to self-identify a potential bot infection.

8. End-User Notification of Potential Bot Infection

a. Overview:

Many end-users are unaware that their devices are infected and operating as bots. As a result, those users and their data remain at risk, and the bots can remain active indefinitely. ISPs should leverage detection efforts described in Section 7 to make customers aware of active infections.

Notifications should be designed to help mitigate bots and the harm they cause. Notifications may include information on what a bot is, means of infection, that bots may have no visible symptoms, and what the notification means. Notifications may also contain or identify other resources such as tools, guides, and services that facilitate infection prevention, verification, and mitigation*. They may also provide information on any specific bot(s) detected.

End-user notification may take many different forms. It may be performed directly by the ISP or by third parties on behalf of the ISP. ISPs may directly alert end-users or provide mechanisms that allow end-users to request and receive information on their infection status. Similarly, ISPs might enter into arrangements to enable notifications to be delivered to end-users by other ecosystem participants with whom the end-user has a relationship, such as a provider of an Internet application or service.

The ISP should consider mechanisms that ensure that the customer is easily able to authenticate the notifications as genuine and that such notifications will be difficult to spoof.

Where feasible, the ISP may wish to track receipt of notifications. This may help the ISP better understand the effectiveness of various notification mechanisms.

Each ISP will need to evaluate different notification methods in order to find one best suited for the particular ISP and the particular bot threat. The notification method chosen may need to integrate with existing business processes and existing network

* Definition found in Glossary

infrastructure. Research and analysis may be required to develop and maintain appropriate notification systems and policies.

b. Recommended Action:

Provide communication of a suspected bot infection to the customer or help enable customers to determine if they are potentially infected by bots. Many notification methods are outlined in references in Appendix 2; however, other methods may be used.

9. Bot Remediation

a. Overview

Bot mitigation and remediation is the ultimate goal of any bot infection notification program and is ultimately the responsibility of the end-user. Notification alone may be sufficient for technical users but the majority of users usually require some form of assistance in removing bot malware from their infected devices. Remediation, however, can be difficult, and may involve other complex functions such as isolating the source of the infection among many devices sharing an Internet connection; backing up all data and system software ahead of time in a way that preserves the end-users' ability to recover (but doesn't back up the infected files or programs as well); and ensuring that the end-user has source disks and other materials from which to reconstruct their device image if required during the remediation process.

It is understood that some ISPs may not have the resources to provide this level of service, nor be able to support such activities free of charge or even for a fee. In many cases, end-users may need to be referred to providers of professional computer support services to fully remediate their machines. ISP notifications may wish to anticipate this fact and suggest that customers seek third party assistance to avoid frustrating end-users with limited-service help-desks or support lines that aren't capable or equipped to fully address the remediation issues.

b. Recommended Action:

1. Bots are designed to be stealthy and difficult to remove. As part of the notification, ISPs should offer guidance, as described above. This may include links to a variety of publically available online and third party sources of information, software, and tools. It might also include links to professional services. These need not be offered by the ISP itself but may be offered by third parties.
2. An ISP may provide remediation tools to the end-user, either during or after the notification process. However, the ISP should not mandate that the end-user run remediation tools. If the ISP provides tools to the end-user, the end-user should be allowed to exit the process without running any suggested tools or procedures.
3. As part of the notification process, ISPs may wish to include guidance (depending on the nature of the bot in question) that settings on customer owned network equipment such as home gateways and routers may have

been altered and should be restored to a secure state, depending on the nature of the bot infection.

c. Guidelines:

1. Bot removal tools and services must respect user privacy.
2. Possible infection remediation methods are outlined in the CSRIC II WG 8 best practices and in the bot remediation IETF RFC which are referenced in Appendix 2.

10. ISP Collaboration

a. Overview:

Bot mitigation and management are activities in which ISPs, search providers, end-users, IT departments, hosting companies, blog providers, security vendors, researchers, government, financial services companies, cloud service providers, and other parties all have roles. With multi-stakeholder input and collaboration, the results will exceed those possible through independent actions alone. ISPs' participation in this code, along with complementary and collaborative approaches taken by other segments of the Internet ecosystem, can be expected to drive substantial mitigation of the threat posed by botnets.

b. Recommended Action:

Code participation requires collaboration within ISP, industry, or broader fora through collaborative activities, of which the following are examples:

1. Sharing detection, notification, or mitigation methods planned for or deployed in ISP networks, and where practical an evaluation of their effectiveness.
2. Sharing of intelligence or operational attack data that may be useful in bot prevention, defense, or remediation.
3. Identification of key data or technical resources that are needed from systems or actors beyond the ISP network.
4. Participation in definition, development, or operation of integrated defense strategies or systems which extend beyond the boundaries of the ISP network.
5. Other collaboration activities involving the sharing of information with parties outside the ISP or data with systems outside of the ISP network.

All information sharing between ISPs and other involved parties will be performed in accordance with applicable laws including, but not limited to, antitrust and privacy laws.

11. Further Development of this Code

This Code will evolve over time due to the dynamic nature of the bot threat and the experience and assessment of ISPs.

12. Additional Information and Resources

Appendix 1 – Glossary
Appendix 2 – References

Appendix 1 – Glossary:

1. Bot

The following definition draws heavily from "Recommendations for the Remediation of Bots in ISP Networks" (Referenced in Appendix 2):

A malicious (or potentially malicious) "bot" (derived from the word "robot", hereafter simply referred to as a "bot") refers to a program that is installed on a system in order to enable that system to automatically (or semi-automatically) perform a task or set of tasks typically under the command and control of a remote administrator (often referred to as a "bot master" or "bot herder.")

Computer systems and other end-user devices that have been "botted" are also often known as "zombies".

Malicious bots are normally installed surreptitiously, without the user's consent, or without the user's full understanding of what the user's system might do once the bot has been installed.

Bots are often used to send unwanted electronic email ("spam"), to reconnoiter or attack other systems, to eavesdrop upon network traffic, or to host illegal content such as pirated software, child exploitation materials, etc.

Many jurisdictions consider the involuntary infection of end-user hosts to be an example of an unlawful computer intrusion.

2. Botnet

Botnets are networks of Internet-connected end-user computing devices infected with bot malware, which are remotely controlled by third parties for nefarious purposes. A botnet is under the control of a given "botherder" or "botmaster." A botnet might have just a handful of botted hosts, or millions.

3. Customer (or "Direct Customer")

The party contracting with an ISP for service. Distinguish the "customer" from an "authorized user:" for example, a coffee shop might purchase Internet service from an ISP. The coffee shop would be the ISP's customer. The coffee shop might elect to offer free use of its connection (if permitted by the ISP's Acceptable Use Policy, or AUP) to those who buy coffee from it -- coffee buyers would then be authorized users of the connection purchased by the coffee shop, but not the ISP's direct customer.

4. Detection

Detection is the process whereby a service provider or end-user comes to be aware that a particular system or device has been infected with malicious software. A service provider may detect that a system has become infected many different ways, including as a result of receiving complaints from third parties about spam, network scanning, or attacks that have been sourced from that system. End-users may detect system infections through software tools or other means.

5. Ecosystem

This term is often used to describe the interrelationship of various Internet participants—the hardware manufacturers, software developers, ISPs, and providers of various Internet content, applications, and services that make the Internet work and be useful for end-users.

The internet ecosystem includes operating system vendors, end-user focused organizations, providers of Internet content, applications, and services, ISPs, search providers, end-users, IT departments, hosting companies, blog providers, security vendors, researchers, government, financial services companies, and other parties.

The so-called "underground economy" is also often described as an "ecosystem," with multiple participants filling diverse specialized roles. For example, some participants may specialize in writing malware, while others may "harvest" email addresses from web pages and mailing lists, while still others may specialize in distributing malware to those harvested email addresses. The malware ecosystem will also normally include the population of targeted potential victims, and law enforcement agencies working to combat cybercrime.

6. End-user

End-User: In a computing and networking context, the end-user is the person who ultimately makes authorized use of a product or service.

The end-user may often not be the same as the person who may have purchased the product or service. For example, a coffee shop owner may purchase connectivity for use by his or her customers; in that scenario, the coffee shop customers, and not the coffee shop owner, represent the actual "end-users," even though they did not directly contract with an ISP for the connectivity they're using.

A party, such as a hacker/cracker who makes use of a product or service without the authorization of the purchaser, would normally be considered a cyber intruder and not an "end-user" per se.

7. ISP

An Internet Service Provider (ISP) is a company that provides retail access to the Internet for members of the public, or for businesses and other organizations. Those connections may be via cable, DSL, satellite, wireless, dialup, or other technologies. ISPs are sometimes also known as "access providers."

An enterprise that provides access to the Internet solely for its own employees would not normally be considered to be an ISP. Likewise, a network carrier that only provides wholesale access to the Internet for other ISPs would normally be considered to be a network service provider (NSP), rather than an ISP.

8. Malware

"Malware" is short for "malicious software."

Malicious bots are one type of malware. Other forms of malware include categories of software known as viruses, Trojan horses, worms, rootkits, crimeware, keystroke loggers, dialers, spyware, adware, etc. The factors that distinguish those different types of malware are less important than an understanding of why malware may be viewed as "malicious."

Malware often violates one or more of the following fundamental principles:

- (a) Consent: Malware may be installed even though the user did not knowingly ask for that to happen.
- (b) Honesty: Malware may pretend to do one thing, while actually doing something completely different.
- (c) Privacy-Respectfulness: Malware may violate a user's privacy, perhaps capturing user passwords or credit card information.
- (d) Non-Intrusiveness: Malware may annoy users by popping up advertisements, changing web browser's home page, making systems slow or unstable and prone to crash, or interfering with already installed-security software.
- (e) Harmlessness: Malware may be software that hurts users (such as software that damages our system, sends spam emails, or disables security software).
- (f) Respect for User Management: If the user attempts to remove the software, it may reinstall itself or otherwise override user preferences.

It all adds up to "software users just don't want."

Users may unknowingly install malware by opening a tainted attachment received by email, or by visiting a web page that has malicious content. Systems may also be directly infected by a remote attacker as a result of the attackers targeting a known vulnerability that may be remotely exploitable, or by the user mounting an infected CD, DVD, or thumb drive.

9. Mitigation

Mitigation is the process of managing or controlling the effects associated with a bot. For example, if a system is infected with a spam bot, and is spewing unwanted commercial email, mitigation may consist of filtering the spam that is being emitted from that device.

Note that mitigation typically does not involve fixing the underlying condition (that would be "remediation"); mitigation just manages the symptoms associated with a condition.

10. Notification

Notification is a process whereby ISPs communicate with their end-users regarding the possible infection of the end-user's device by bot malware or how a subscriber can prevent or identify such an infection. Notification may also entail a process whereby end-users are directed to tools that will enable self-discovery of bot infections. Notification can take different forms, including direct notification by the ISP to the end-user, or indirect notification through available self-discovery tools or a third party. Notification may be done via multiple potential channels, including (but not limited to) e-mail, postal mail, a phone call, in-browser notification, web-based self-discovery tool, or SMS message.

11. Prevention

Prevention is the process of hardening a system or service so that it is less vulnerable to compromise and exploitation. For example, on many systems, prevention may involve:

- Patching the operating system and all applications with available security fixes
- Installing or enabling a firewall
- Using anti-virus software
- Making sure the system is regularly backed up
- Using strong passwords
- Disabling all unneeded network services
- Encouraging users to safely use internet services (e.g., e-mail, web browsing, etc.)

12. Remediation

Remediation is the process that an end-user goes through to clean up a botted computer so that it is no longer infected. In easy cases this may involve installing and running an anti-virus product. In more difficult cases, remediation may involve more substantial intervention up to "nuking and paving" the system -- formatting it and reinstalling it from scratch, or at least from the last known-clean backup. Once the system is clean, or has been reinstalled, it will then normally be hardened to protect it from reinfection.

13. Spam

Unwanted and unrequested e-mail, often commercial in nature, normally sent to a large number of recipients in substantially identical form. Spam is often sent by "affiliates" who are paid by the person running the affiliate program when recipients purchase the spamvertised product.

Appendix 2 – References

1. Recommendations on how to manage the effects of computers infected with malicious bots: "Recommendations for the Remediation of Bots in ISP Networks"

<http://tools.ietf.org/rfc/rfc6561.txt>

2. CSRIC II Working Group 8 - ISP Network Protection Best Practices

http://transition.fcc.gov/pshs/docs/csric/CSRIC_WG8_FINAL_REPORT_ISP_NETWORK_PROTECTION_20101213.pdf

3. icode – Australian Internet Industry Code of Practice addressing cyber security

<http://iia.net.au/images/resources/pdf/icode-v1.pdf>

4. Japan Cyber Clean Center – Anti-Botnet Project

https://www.ccc.go.jp/en_index.html

5. German Anti-Botnet Advisory Centre – Anti-Botnet Project

<https://www.botfrei.de/en/>

6. Japan Computer Emergency Response Team (CERT)

<http://www.jpCERT.or.jp/english/>

7. US CERT - Understanding Hidden Threats: Rootkits and Botnets

<http://www.us-cert.gov/cas/tips/ST06-001.html>

8. Alliance for Telecommunications Industry Solutions (ATIS)

<http://www.atis.org/>

9. Department of Homeland Security

http://www.dhs.gov/files/programs/gc_1158611596104.shtm

10. Department of Homeland Security - United States Computer Emergency Readiness Team (US-CERT)

<http://www.us-cert.gov/>

11. International Telecommunication Union Botnet Mitigation Toolkit

<http://www.itu.int/ITU-D/cyb/cybersecurity/projects/botnet.html>

12. U.S. Commerce Department's National Institute of Standards and Technology (NIST)

<http://www.nist.gov/index.html>

The Communications Security, Reliability and Interoperability Council
Final Report

Working Group 7
March, 2012

13. Department of Commerce/Department of Homeland Security Request for Information - Models To Advance Voluntary Corporate Notification to Consumers Regarding the Illicit Use of Computer Equipment by Botnets and Related Malware

<http://www.gpo.gov/fdsys/pkg/FR-2011-09-21/pdf/2011-24180.pdf>

14. Comments Received in Response to Department of Commerce/Department of Homeland Security Request for Information - Models To Advance Voluntary Corporate Notification to Consumers Regarding the Illicit Use of Computer Equipment by Botnets and Related Malware

<http://www.nist.gov/itl/botnetcomments.cfm>

15. Messaging Anti-Abuse Working Group (MAAWG.org) - Code of Conduct

<http://www.maawg.org/sites/maawg/files/news/CodeofConduct.pdf>

16. M3AAWG Collection of Best Practices for ISPs and Network Operators

<http://www.maawg.org/published-documents>

17. National Vulnerability Database – National Institute of Standards and Technology

<http://nvd.nist.gov/>

18. Internet Storm Center

<http://isc.sans.edu/index.html>

19. Shadowserver Foundation

<http://shadowserver.org>

20. Spamhaus Policy Block List

<http://www.spamhaus.org/pbl/>

21. Composite Blocking List

<http://cbl.abuseat.org>

22. OnGuard Online

<http://www.onguardonline.gov/default.aspx>

23. IETF BCP38 Network Ingress Filtering

<http://tools.ietf.org/html/bcp38>

Appendix 3 – Barrier Guide



March 2013

Appendix 3 Guide to Barriers to Code Participation

U.S. Anti-Bot Code of Conduct (ABC)
for Internet Services Providers (ISPs)

WORKING GROUP 7 - Botnet Remediation

Introduction

The CSRIC Working Group 7 created the Guide to Barriers to Code Participation (“the Guide”) to provide ISPs with guidance on Code implementation activities based on the experience of the working group and other current information. It is hoped that this information can guide ISP efforts at Code participation, and help answer some of the more common questions ISPs are likely to face.

The guidance is grouped according to the five primary areas for participation under the Code.

- I. End-User Education
- II. Detection
- III. End-User Notification
- IV. Remediation
- V. Collaboration

In addition, each of the five areas is further broken down into sub-parts corresponding with the recommendations contained in the Code.

Finally, the commentary in each of these areas is further sub-divided to try to distinguish between the various types of barriers that ISPs are thought likely to face.

- I. Technology Barriers
- II. Consumer and Market Barriers
- III. Operational Barriers
- IV. Financial Barriers
- V. Legal, Regulatory, or Policy Barriers

Recommendations, Considerations, and Guidance

In compiling the information contained in this Guide, working group participants were guided by the following questions:

1. What are the minimum technical/systems required to implement the recommendation (one-time, on-going)?
2. What is the degree of systems integration that is required to implement the recommendation (one-time, on-going)?
3. What technical resources must be available within an ISP to implement the recommendation (one-time, on-going)?

4. What technical capabilities and resources are available from third parties to implement the recommendation (private, not for profit, non-profit, other)?
5. What is the desired consumer reaction to the recommendation when implemented? How are consumers expected to react? How have they reacted when this recommendation was implemented?
6. Has the system been reviewed from a human factor perspective? If yes, what has been learned?
7. To what extent is implementation of this recommendation transparent to the consumer? Does that matter? Why/Why not?
8. Is there any evidence that consumers who are impacted by a recommendation change their behavior over time? Is there any data on recidivism?
9. What is the level of customer support/care that is associated with this recommendation?
10. What are the dominant concerns that consumers or consumer representatives express?
11. What survey evidence is publicly available to validate consumer reaction? What anecdotal information can be shared?
12. What are the minimum operational process changes that are required to implement this facility (one-time, on-going)?
13. What existing operational processes are impacted by implementing the recommendation (one-time, on-going)?
14. What is the degree of operational integration (e.g., inter-intradepartmental) that is required when implementing the recommendation (one-time, ongoing)? For example, has the system(s) been integrated with customer call center systems?
15. What measurements are required to ensure that this recommendation is producing desired results?
16. What legal, regulatory, or policy considerations (e.g., privacy, anti-trust, and tort liabilities) need to be addressed to implement this recommendation?
17. Are we aware of any international rules or policies that are implicated by this recommendation?
18. What benefits have been quantified associated with implementing this particular recommendation? Is this information publicly available? Can it be obtained through other mechanisms? If yes, which ones?
19. What are the one-time and recurring costs associated with the design, implementation and management of new technologies, systems, processes, and policies associated with implementing this recommendation (in-house/out-sourced)? How do economies of scale and scope impact these costs?
20. What are the cost savings associated with implementing this recommendation? How does the recommendation affect customer support, service performance, and customer loyalty?
21. Is it possible to measure the return-on-investment (ROI) from implementing this recommendation? If not, how should ISP technical managers make the case to the CFO?

22. What metrics/measurements are available to measure effectiveness of implementing this recommendation? What information can be, or is being shared?
23. What is the level of customer support/care that is associated with this recommendation and have care costs been quantified? What factors account for cost variability?

Not all of the above questions were able to be answered for each of the recommendations below. In some cases, more research or information is required. However, it is expected that the information provided below will prove useful to ISPs as they commence their Code participation efforts.

END-USER EDUCATION

Recommendation # 1 PREVENTION: ISPs should make available information about the prevention of bot infections and related issues. At a minimum, such information should include: 1) how and why end users must keep their software updated for devices with readily available software updates; 2) the importance of using effective and current security software from a reputable vendor; 3) the importance of backing up user data and software; and 4) basic end-user actions to minimize bot infections when using the Internet.

Technology Considerations

A minimum capability in this area would include a simple webpage with links to general resources such as www.onguardonline.gov and www.staysafeonline.org. A more sophisticated educational site might provide software update screenshots, links to anti-virus services, and information on application update tools. These links should be reviewed, validated, and updated as necessary.

ISPs may wish to publicize the availability of such a page via a link from the ISP's homepage, bill inserts, and other subscriber communications both electronic and printed. Cross-linking from the ISP's Privacy Policy and security pages may also be done.

Typically, systems integration is not required as part of these efforts. ISPs typically utilize existing web platforms and in-house expertise in web site creation and maintenance to build and maintain such a website. To the extent the ISP does not have such resources in-house, website creation, maintenance, and hosting expertise is typically available from a range of third-party vendors.

Consumer/Market Considerations

The intent of these activities is to inform customers of resources and educational materials so that they can take basic, effective steps to secure their systems and engage in safe online

behavior. It is expected that some portion of users will understand and take prescribed action. How consumers have reacted to ISP efforts in this area so far is not well known.

ISPs should solicit input from their internal usability specialists or retain third parties for simple assessments to ensure that the materials prepared by the ISP are easily understandable to the intended audience.

Some ISPs include references to their own paid anti-virus and security products as part of these activities. Typically, this is done in conjunction with information on other similar resources to avoid tainting the security-focused message with an air of marketing activity.

There is little evidence to-date that consumers change their behavior in response to such information, but it is expected that better-informed end users are generally more knowledgeable on topics such as simple social engineering exploits, and may take steps to reduce their vulnerabilities. It is suggested that ISPs at a minimum collect page views or click-through traffic for resources posted either on their own or third-party sites so they have some relative measure of the extent to which end users are visiting such resources.

Typically, there is little to no impact on ISP customer care centers in conjunction with the provision of basic educational information.

Operational Considerations

ISPs typically require little in the way of new operational processes to implement this recommendation. Some new processes are required to ensure the web pages are kept current and up-to-date. Support teams need to be advised of site changes and new resources.

In addition, collaboration may be required between customer care, marketing, and legal, particularly if the resource pages are to include links to the ISP's paid security products and services.

Legal/Regulatory/ Policy Considerations

As with any information posted online, legal requirements may vary from state-to-state, but in general, ISPs find that standard disclaimers and caveats suffice. That said, ISPs are advised to consult with legal counsel and have their web content reviewed by their attorneys. Attention should be paid to legal issues associated with linking to third-party websites, re-using third-party materials, and collection of web traffic statistics in light of the ISP's privacy policy.

Because the focus of the Code is on U.S. ISPs, there are not expected to be any international legal or policy implications from providing educational information online focused at a U.S. subscriber base.

Financial Considerations

The benefits of providing educational information online have been hard to quantify. ISPs typically justify such activities through qualitative measures such as customer good-will. Some ISPs are able to associate their overall Code-participation activities with a variety of quantifiable benefits such as reduced churn (customer retention) and reduced network traffic (and therefore reduced bandwidth costs due to the elimination of unwanted traffic associated with botnets). However, these benefits are usually quantified at the level of the overall program, not simply the educational activities.

Implementing educational materials will involve minimal recurring costs once established. There is an argument to be made for end-user education centers being shared or centralized (such as the examples cited at the beginning of the discussion of Technology Considerations). This can substantially reduce ISP costs associated with developing end-user educational content, and can ensure that end users receive consistent messages across ISPs. That said, the cost of implementation typically has been minimal.

Recommendation # 2 END-USER REMEDIATION: Along with information on prevention, ISPs should make available (e.g., via ISP publications, third-party publications or web links) information on how end users can generally remediate bot infections. In this area, it is expected that ISPs will be able to accomplish this goal by linking to existing, publicly-available sources of such information or by creating new sources of such information, either individually or in conjunction with others.

In connection with an ISP's end-user notification activities, ISPs should include in such notices or otherwise, information on where the recipient might turn for additional information and assistance. Such information might include links to publicly-available online information, security tools, or suggestions to seek assistance of a computer professional. Additional subjects and references that an ISP might wish to include are:

- **Risks to the end user and Internet community from using a device that is believed to be infected,**
- **How to identify and remove common forms of bot infections,**
- **Publicly-available tools or services (free or paid) to assist in the detection and removal of bot infections, and guidance on where to find additional (free or paid) assistance.**

Technology Considerations

General educational information on remediation can often be addressed in a manner similar to the forgoing guidance on educational information regarding prevention. Indeed, the twin topics of prevention and remediation are addressed on many third-party sites, and can be incorporated into an ISP's overall general effort at end-user education. The information on remediation provided at the education stage typically stands in contrast to the more-specific guidance that is

provided to end users once an infection has been identified insofar as remediation information at the education stage is provided at a general overall high level. There typically are no new technology considerations associated with providing educational information on remediation that are not otherwise present with providing educational information on prevention.

Consumer/Market Considerations

General information on remediation presented as part of the educational process may help consumers fully appreciate the consequences of an infection, and possibly incent them to be more careful in their online habits before they are infected.

Operational Considerations

In general, providing educational information on remediation requires the same sort of operational processes as does providing educational information on prevention.

Legal/Regulatory/ Policy Considerations

In general, providing educational information on remediation raises the same sort of legal issues as does providing educational information on prevention.

Financial Considerations

In general, providing educational information on remediation raises the same sort of financial issues as does providing educational information on prevention.

DETECTION

Recommendation # 3: ISPs can find out about malicious activity and bot compromised end-user devices by receiving information from external entities, particularly those designed to aid with the overall understanding and real-time dissemination of bot related data.

Technology Considerations

The technologies required to receive information on botnet infections can vary greatly. At the simplest level, ISPs need to be able to download or receive electronic lists of infected IP addresses. There are a number of third-party organizations that purport to provide lists of bot-infected IP addresses. ISP handling of a one-time bot infection event can be handled through manual methods, such as text files and spreadsheets. Correlation of subscriber information can be done manually or with temporary scripts. This method can provide ISPs clear insights into the systems that must be developed to support an on-going program.

On-going programs that routinely intake and process data in text files requires minimal to moderate resources, depending on existing ISP systems. Moderate resources are required to develop a system to track subscriber information and track notifications sent to subscribers, depending on the notification method. If an ISP has already developed systems to correlate subscriber information from an IP address and timestamp for other purposes, the level of effort is much less in building a bot detection and notification system.

ISPs may wish to implement automated notification systems and databases that track incoming, third- party lists, correlate subscriber information for notifications, and track notifications sent to the subscribers. It is highly recommended that ISPs have someone review third-party data and subscriber notifications until the quality of the third-party data feeds have been verified by the ISP. False positive notifications can be detrimental to a bot notification program both from a customer and ISP perspective.

Integration into the subscriber IP tracking system is likely to be required for notification. Timestamps are especially important to ensure notifications are sent to the proper subscriber where dynamic IP addresses are assigned. Such integration is also often leveraged to provide customer support personnel with visibility into which subscribers are notified.

There are several commercial and non-profit organizations that provide high quality bot infection data feeds to ISPs. These data feeds do not require equipment to be deployed in the ISP networks although deploying some equipment can be helpful. These entities employ a number of techniques to watch for infected subscribers including dark IP monitoring, honeypots, Domain Name System (DNS) monitoring, email behavioral monitoring, botnet command and control monitoring, and other techniques to provide accurate, high quality data. These organizations typically segment the data by ISP Autonomous System (AS) numbers or IP blocks and send the ISP infection data for only their customers.

There are no readily available non-profit solutions in this space, although an ISP with an especially technically focused staff can build a highly successful solution from open-source software components, if they have the resources to do the development and integration work themselves.

As the ISPs move from the current IPv4 world over to a dual IPv4 & IPv6 infrastructure, and then on to IPv6 only, the tools developed will need to cover both IPv4 and IPv6 detection and correlation.

Consumer/Market Considerations

It is important that ISP customers have a clear understanding of the botnet problem, the risk bot infections pose to their personal or company information, and the safeguards ISPs put in place to

protect the privacy of their customers' information. This is typically addressed through the education process.

Customers can be very concerned about ISPs monitoring their information. When some customers become aware of the program, they may express concerns around the privacy of their information. ISPs should develop clear talking points for the media around their programs including the FCC recommendations for implementing such programs and how such programs help protect the security and privacy of the customers' information. This information can be posted on the ISP's website alongside general educational materials so that customers are aware of the ISP's Code participation activities and of what to do when notified about an infected device.

ISPs should review their Acceptable Use Policy (AUP) and other customer contract information, to make sure that any proposed bot detection program is covered by such a policy.

ISP subscribers are generally unaware of third-party monitoring systems. With proper media and customer talking points, concerned customers can be educated concerning the malicious behavior that the infections are performing and the impact the infection has on other Internet users.

Anecdotal evidence shows that consumers are more aware of Internet security issues and take steps to better protect themselves when clear, understandable guidance is provided by the ISP or third parties. More study is needed in this area.

Detection typically does not trigger customer care costs—customer care is more often required at the notification stages, although the degree of customer care can be mitigated by clear communication and provision of adequate guidance on remediation.

Case studies are underway to study the effectiveness of bot infection notification and remediation programs and techniques in general. Further study is needed to understand consumer reactions to these programs.

Operational Considerations

ISPs must track subscriber IP assignment information and track that information over time. This capability is typically already in place to support ISP subpoena response activities.

In addition, ISPs must validate the quality of third-party infection information. For a program to be successful, the data must have a low false positive rate. False notifications may dramatically undermine an ISP's program, generating high call volumes and loss of consumer confidence in the program.

It is important for an ISP to understand that it can be very difficult for consumers to understand the infection notification, find the infected device, and remediate the problem. Factors that compound the problem of consumers successfully addressing infections include:

- Transient end users who temporarily use the subscribers' systems including friends, relatives, and school mates.
- Unsecured Wi-Fi access points where infected neighbors use the customer's Wi-Fi without their knowledge or permission.
- The large number of Internet connected devices in the home including computers, smartphones, VoIP phones, televisions, DVD players, streaming media devices, home security systems, routers, modems, Heating, Ventilation, and Air Conditioning (HVAC) systems, and appliances. Most of these devices are not managed or even known by the ISP.
- The inability of most ISPs to see beyond the broadband modem or router to identify the infected device in the customers home or business.

Given this complex environment, false positive notifications can dramatically consume ISP resources trying to help a subscriber remediate a non-existent problem. Detection must be handled accordingly.

Legal/Regulatory/ Policy Considerations

ISPs should consult with legal counsel before implementing detection programs. In-house programs need to comply with existing law and privacy policies. Third-party solutions often require vendor contracts with terms of service. ISPs who use third-party infection data should understand the limitations that the third-party places on the use of the data.

Financial Considerations

Using third-party subscriber infection feeds can be a very cost-effective method of obtaining customer infection information when an ISP develops their program. High quality feeds are available from non-profit organizations and some federal government agencies (generally in conjunction with law enforcement actions).

ISPs will incur costs to develop and maintain the subscriber infection database and correlation system. Typical costs include server equipment, database and system development, and system maintenance costs.

DETECTION

Recommendation # 4: Deploy capabilities within their networks that aid in identifying potential bot infections.

Technology Considerations

In-house detection tools can use a variety of methods for identifying infected customers, ranging from relatively lightweight methods such as NetFlow and DNS traffic analysis up to a complete Deep Packet Inspection (DPI) solution. The technologies used to implement such solutions can vary widely, based on the form of the ISPs network, the specifics of their infrastructure architecture, and the contract/AUP covering the relationship with their customer(s). Regardless of the technologies chosen for in-house detection, the resulting data will need to flow through a process that correlates IP address to customer in preparation for notification activities.

Ideally, any bot detection program should be completely transparent to the customer with the only customer awareness coming from notification and remediation work. Detection systems should be developed such that a failure or outage of the detection system is completely non-interfering with the customer traffic and internet access.

Bots are a constantly changing problem, evolving quickly by active criminal elements to work around whatever blocks are put in their way. As such, methods that are found to be highly effective today may lose their value over time as bot writers move to alternative designs. As such, systems implemented by ISPs will need to be constantly re-evaluated to ensure their accuracy and success.

Also, as the ISP moves from the current IPv4 world over to a dual IPv4 & IPv6 infrastructure, and then on to IPv6 only, the tools developed will need to cover both IPv4 and IPv6 detection and correlation.

As with third-party notifications above, the primary integration work will come with being able to correlate an infection to a customer given the observed IP address and time of detection.

In-house detection systems can be relatively hands-on or hands-off, based on the desires of the ISP. Systems that the ISP pulls together from a variety of different solutions will obviously require greater technical oversight and integration compared to a “black box” purchased from a vendor.

Otherwise, the requirements here match-up quite closely with the requirements for a third-party based solution, namely, a system for correlating against customer data.

There are several vendors that provide relatively turn-key solutions for in-house detection. There are no readily available non-profit solutions in this space, although an ISP with an especially technically focused staff can build a highly successful solution from open-source software components, if they are willing to do the development and integration work themselves.

Consumer/Market Considerations

As stated above in the third-party data feed section, customer education on the risks of bot infections is important for a variety of reasons. Any in-house detection system should be

developed such that it has no impact on customer traffic, even if the detection system performs improperly or fails.

Care must also be taken to ensure that customers do not misconstrue network-based bot detection as an invasion of privacy.

There should be no customer support implications for an in-house detection system. Such a system will require standard technical and operational support.

Operational Considerations

As stated above, there are operational implications to the system required to map IP address to customer. Since this is required to provide support for subpoenas and other legal demands, this is unlikely to be problematic. In addition, depending on the method selected for in-house detection, there may be operation implications of the running and maintaining of such a system.

At a minimum, any in-house detection system will need to be monitored for operational status, and standard hardware maintenance done on any equipment deployed to implement this solution.

Legal/Regulatory/ Policy Considerations

Care must be taken to ensure that a subscriber's privacy is maintained. Because network-based detection tools may need to inspect customer traffic, they must be designed in such a way that they do not store Personally Identifiable Information (PII) or violate any other applicable laws. ISPs should consult with legal counsel prior to implementing a detection program to ensure that such a program is implemented in accordance with applicable law.

Financial Considerations

The financial considerations for an ISP-developed detection platform can vary greatly depending on the technology chosen (NetFlow, DNS, other), the hardware and software platform selected, the performance parameters required, the network architecture and volume of traffic, and other factors.

DETECTION

Recommendation # 5: Direct customers to tools, a web portal, or other resources that enable customers to self-identify a potential bot infection.

Technology Considerations

This approach involves maintaining lists of infected IP addresses and offering access to that information as a service that enables subscribers to periodically check to ensure their device is not infected. ISPs would need to develop database and portal systems to provide this capability.

ISPs would need to authenticate the subscriber and then provide information with respect to whether or not that subscriber's account has been seen as being engaged in botnet activity.

Consumer/Market Considerations

Information would need to be developed for consumers to clearly indicate how the service is provided and what the results mean.

Operational Considerations

The operational considerations are similar to botnet detection through other means, but also include operational capabilities to maintain the service through which the infection information is communicated to subscribers.

Legal/Regulatory/ Policy Considerations

ISPs are advised to consult with legal counsel. Many of the issues related to detection generally may well be applicable to this type of service.

Financial Considerations

Relevant financial considerations are similar to those present with other detection methods.

END-USER NOTIFICATION

Recommendation # 6a: Provide communications of a suspected bot infection to the customer to help enable customers to determine if they are potentially infected by bots. Many notification methods are outlined in references in Code, Appendix 2; however other methods may be used.

Recommendation # 6b: Enable customers to determine if they are potentially infected by bots. Many notification methods are outlined in references in Code, Appendix 2; however other methods may be used.

Technology Considerations

There are a wide range of communications mechanisms, all with varying degrees of cost, flexibility, and effectiveness. ISPs participating in Code-related activities have taken various approaches to notification, including voicemail messages, Short Message Service (SMS), postal

mail, email, account messages, network walled-gardens/browser redirection, and in-browser messaging.

Some ISPs have found success in inducing end-user remediation behavior through email and other such communications. Other ISPs have found better success once notification efforts have escalated to more persistent/intrusive mechanisms, such as in-browser notifications, browser redirection, and network walled-gardens.

However, because of multiple users in a household, technologies that target the web browser may not reach the primary account holder. Rather, the notification may be seen and dismissed by another member of the household or by a neighbor using the subscriber's unsecured wireless access point. In addition, care must be taken not to display personally identifiable information about the subscriber in such a notice without requiring authentication.

ISPs recommended initially taking a graduated approach that progresses towards more intrusive actions (e.g., browser redirection, network walled-garden, suspension of service) only after repeated notification efforts have been ignored.

At the initial stages of such escalation, typically, any end user is permitted to acknowledge the notice or "click-out" of the walled-garden or redirection technology. These are often referred to as "leaky" walled gardens. At later stages (where repeated notices have been delivered), some ISPs have implemented procedures that require the subscriber-of-record to call-in to the ISP, authenticate themselves, and discuss their remediation plans/efforts.

Implementing a graduated program requires an ability on the part of the ISPs to track notifications by subscriber. This increases the costs and complexity of a notification program. Some ISPs that have implemented such tracking, however, have found it useful as a means of enabling personnel in various business functions (abuse, customer support) to have access to a common repository of such communications to facilitate subsequent conversations with subscribers.

Implementing a notification capability involves being able to identify the subscriber account information associated with botnet activity associated with a particular IP address at a particular date and time. As noted earlier, the capability to link a subscriber with an IP address by means of a time stamp is already likely to be present in order to meet legal demands such as subpoenas. Once this association is made, notification can be made using a variety of notification methods with varying degrees of automation.

The degree of systems integration that is required may vary substantially. The more extensive the notification program (such as notifying on greater numbers of bots, notifying increasing numbers of subscribers, notifying in a quicker and more automated fashion, notifying using in-browser notifications, or walled-garden implementations), the greater the degree of systems integration likely required.

ISPs likely have the existing technical capability to link a subscriber account to a bot infection, and then notify that subscriber via email or other means, but the level of resources required to implement a substantial and ongoing notification capability may not exist and may require additional resources.

ISPs likely have existing systems to enable subscriber identification associated with IP address utilization, and separately, likely have systems that enable communication with their subscribers. However, as noted above, the degree of systems integration that may be required in connection with particular notification efforts may be substantial.

Traditional IT consulting services are available from a wide range of vendors, but may require significant costs depending on the level of systems integration desired.

Consumer/Market Considerations

Care should be taken to craft a notification that is understandable and effective. The moment of notification is an opportunity to deliver a short, precise message that elicits the desired response on the part of the end user (obtaining suggested resources for, and implementing, system remediation).

ISPs participating in Code-related activities have taken various approaches to the content of notifications. Most notifications provide a brief summary of the issue, with references to where end users can find additional information. These references may be to the sources of information on end-user remediation maintained on the ISP's website or on third-party websites.

Some ISPs have found success by paying close attention to the language used in the notification text—for example, although industry groups often speak in terms of “malware” and “botnets” some ISPs have found greater success notifying end users of a “computer virus” infecting their system.

Some industry participants have called for notifications to pay careful attention to varying levels of end-user sophistication, language, dialect, and ability. Attempting to tailor notification efforts in this manner, however, introduces additional complexity and cost into such programs. ISPs participating in notification activities today typically do not tailor notifications in these manners. This is an area where third-party commercial solutions may be created as notification efforts become more widespread. In addition, end users typically have access to third-party professional services for remediation, so the point/effect of the notification can be reasonably limited to informing the end user of the need to take action.

Recipients of notices are expected to take steps to identify and remediate their infected devices.

As discussed, above, ISPs typically review their notices from the human factor perspective.

Notification is not transparent. Subscribers must receive and act on the notices. In addition, some ISPs have selected notice methods such as browser-inserts and walled-garden capabilities that are visible within or disrupt the web-browsing experience.

Some ISPs have found that their notification activities trigger relatively few calls to their help desk. If the notice is reasonably crafted and delivered, most subscribers appear to be capable of understanding that they have a computer virus and need to take steps to remove it. In addition, the notice typically directs the subscriber to web-based resources where they can get additional information on remediation. The information on remediation typically does not include suggestions to call the ISP, but directs the end user to self-help resources and third-party paid professional support. However, most ISPs take steps to ensure that customer-care personnel are aware of the ISP's botnet notification activities, including which subscribers have been notified, in order to streamline the customer-care experience should those customers call with questions.

Subscribers typically have questions regarding why they received the notice and what they should do next. Having good, web-based resources on both topics can help minimize calls to the ISP.

Anecdotally, many ISPs have found that subscribers are pleased to have been informed about the situation, once they understand the meaning of the notice.

Operational Considerations

There is a wide range of malicious software, and ISPs need to make decisions regarding when to provide notification. ISPs participating in Code-related activities today have taken various approaches.

Some provide notification only when there is an automated "fix" or program that the end user can be directed to for remediation. The concern is that providing notification without a solution may be counter-productive, if it leads to frustration and hopelessness on the part of the end user, and diminishes the end user's willingness to review and respond to future notifications.

Other ISP's provide notification whenever malicious software is detected by the ISP's technology even if an automated software fix is not yet available, so that end users can change their behavior (e.g., avoid banking online for a period of time) and engage in remediation either on their own or with professional assistance.

Typically, ISPs need to implement new processes to associate botnet infection data with subscriber accounts and then deliver an appropriate notice. The range and complexity of new processes required will vary greatly depending on the approach taken by the ISP.

ISPs have typically found that the new processes primarily impact the Security & Abuse functions, but customer-care is often impacted as well since it is advisable to ensure that

customer-care personnel are aware of the notification efforts and of where to send subscribers calling to seek additional information.

The degree of process change may vary by ISP. To the extent ISPs are not already associating botnet infection information with subscribers and notifying those subscribers, new processes would have to be developed to support that activity.

The degree of operational integration, as with systems integration, will vary greatly based on the extent of the ISPs notification activities. At a minimum, Security & Abuse teams likely need to interface with customer-care groups to ensure both are aware of notification campaigns.

ISPs engaging in notification activities try to track subscriber notifications both to enable more efficient customer-care, and to help determine whether the customer took action. However, more work is required in the area of metrics.

Legal/Regulatory/ Policy Considerations

The scope of ISP liability for activities taken in connection with Code participation may vary, depending on the way in which the ISP implements its program.

All ISPs currently participating in Code-related activities consulted with Legal counsel in connection with implementing their programs, and made suitable enhancements to existing legal documentation such as terms of service, acceptable use policies, privacy policies, etc.

ISPs exploring participation in the Code are encouraged to take an approach that makes sense within their particular network, technology, and operating environment. ISPs may wish to consult with legal counsel regarding their notification program to ensure that measures are taken to limit liability to the ISP, for example, for failing to notify of known issues, for notifying and inducing end users to take remedial measures that cause further damage to or loss of end-user systems or data, and for other aspects of their program.

As noted above, in connection with implementing any of the recommendations, ISPs need to work with legal counsel to address legal and liability issues.

It is not clear that any international rules or policies are implicated by the recommendations.

Financial Considerations

Financial considerations must be taken into account. Any program must be sustainable financially in the sense that the benefits from the program outweigh its costs.

ISPs have considered a variety of financial benefits when implementing their programs, including customer good will, reduced churn, reduction in spam or other malicious traffic on the network and reduction in calls to their help desks from customers with slow, infected machines.

Few ISPs have been able to precisely quantify the financial benefits. More research needs to be done in the area of quantifying benefits of notification.

Many ISPs currently operating anti-bot programs have found Code-participation to generate few help desk calls. If presented with appropriate notification and self-help resources online, these ISPs have found that their end users appear to be able to deal with a “computer virus” on their system without calling their ISP.

ISPs typically incur one-time and recurring costs associated with making a process and system changes to engage in notification. Automation and systems integration has enabled ISPs to reduce the marginal costs associated with ongoing or additional notification activities.

Some ISPs have experienced costs savings in terms of reductions in traffic due to eliminating spam and other malicious traffic from their networks. However, more research is required in this area.

ISPs typically justify the business case for Code participation based on qualitative factors, such as assumptions around customer retention and customer goodwill. It may be possible for some ISPs to quantify savings based on reduced transit costs from upstream providers due to reduction in unwanted network traffic, and reduced calls to ISP help desks from subscribers who experience computer difficulties due to infections that the subscriber is incorrectly attributing to slow performance of the ISPs access service. ISP notification activities may also drive revenue for other ISP services such as managed or subscription-based security services and professional services

Some ISPs have found little to no additional cost associated with customer care in the wake of undertaking notification. Other ISPs have found that notification activities do drive some additional calls to their help desk. Increasing end-user education and awareness, having carefully crafted notification methods, and good online resources can all help guide the subscriber through the remediation process without the need for a call to the help desk. In addition, having ISP customer-care personnel informed of the notification campaigns and the subscribers that were notified can enable customer-care representatives to quickly direct callers to the appropriate resources for assistance, thus minimizing call times.

REMEDICATION

Recommendation # 7: Bots are designed to be stealthy and difficult to remove. As part of the notification, ISPs should offer guidance, as described above. This may include links to a variety of publically available online and third-party sources of information, software, and tools. It might also include links to professional services. These need not be offered by the ISP itself but may be offered by third parties.

Technology Considerations

ISPs can either develop their own remediation information or rely on third-party information. Many times, larger ISPs have the resources to develop their own material in order to better control the subscriber experience. Usually, this is done in coordination with security researchers, anti-virus vendors, and operating system vendors. Small ISPs can leverage third-party information and refer customers to third-party sites to minimize costs.

Minimal resources are required to link to third-party or professional remediation services. Some additional effort is required if the ISP wishes to maintain a legal relationship with a third-party tool vendor. Additionally, an ISP must ensure third-party and internally developed information remains active and relevant.

In-house development requires more extensive resources but provides the ISP with better control of the end-user experience.

ISPs that rely on third-party sites should ensure that links stay accurate and are routinely reviewed.

With in-house development, the ISP with their own technical support, web pages and support organization will require more extensive integration. Additionally, integration with third-party support sites (free and paid) may also be required.

ISPs should try to ensure that any technical advice provided to customers has a reasonable chance of successfully remediating their infection. However, not all ISPs actually conduct tests ahead of time of third-party remediation tools—it would be difficult and cost-prohibitive to do so given the wide variety of end-user devices and configurations.

Generally speaking, new systems will not be required but existing technical support sites and processes will need modifications. Operating system and anti-virus vendors can supply free or commercially provided tools to assist consumers in cleaning many of their devices. Third parties can also provide commercial services to help consumers clean their infected devices. ISPs may also give consideration to operating their own services to assist customers in remediation.

Consumer/Market Considerations

Ideally, the information provided should allow a consumer to either clean their infected device or be pointed to professional resources that can assist the consumer in remediating their problem.

Clear, easy-to-understand instructions are important to minimize support costs to the ISP.

The level of customer support/care that is associated with this recommendation depends greatly on the approach taken by the ISP towards providing remediation services to the customer. If the ISP chooses to rely on third-party products for remediation, the ISP may be able to leverage that third party for support during the remediation process. More in-house managed solutions will

require more in-house support. Anecdotal information suggests that remediation is an area where extensive online information and self-care tools can reduce call volume, but more research in this area is required to validate those assumptions.

Remediation tools are not perfect. In fact, depending on the type of the infection, they may be difficult to use or completely ineffective. This can drive customer complaints and calls. Based on feedback received by ISPs who are currently doing large scale remediation and represented on Working Group 7, the general customer reaction is one of gratitude to the ISP for notifying them of the infection. As for the remediation tools, most customers do not have strong opinions unless the tools fail. 100% effectiveness is the expectation from the consumer, and this is not always met.

Operational Considerations

ISPs should consider the ability of criminal actors to imitate legitimate links and trick consumers into going to illegitimate sites. ISPs should anticipate these scams and take steps to minimize and mitigate them.

Bots constantly change and evolve, and the tools to fight them change frequently as well. ISPs should ensure that guidance given regarding remediation of specific bots is reviewed periodically to be kept current. At a minimum, an ISP providing remediation services should make efforts to verify that the services work, and should ensure that any links to third-party products are accurate and timely.

Depending on how the remediation tools are linked into the existing ISPs infrastructure, changes to tools may require updates to multiple locations as well as updates to customer care scripts.

ISPs need technical resources to ensure that they can identify resources that might be useful to the consumer. Generally speaking, new systems will not be required but existing technical support sites and processes will need modifications.

Legal/Regulatory/ Policy Considerations

ISPs should ensure that they have reviewed all legal requirements for the use of tools. They should consider language in their program around liability of cleaning instructions and the use of cleaning tools or services.

As always, ISPs should follow appropriate state and local laws in the areas where they offer service.

Financial Considerations

Some ISPs have internally documented a drop in support calls from consumers without malware infections versus consumers with malware infections. Due to their nature, bot infections can

appear to the subscriber to be a poor broadband service leading to higher support calls, up-stream bandwidth utilization and higher customer churn.

ISPs should evaluate these factors in their business case as they consider a bot notification and remediation program.

Leveraging third-party cleaning information and referring consumers to third-party providers is the least costly short-term method in end-user remediation options. ISPs must consider customer experience and longer-term support costs when making decisions on their program.

Reduction in call volumes and up-stream bandwidth utilization are cost savings that some ISPs have experienced associated with Code participation.

In terms of return-on-investment (ROI), ISPs can setup tests to measure customer technical support call volumes and lengths associated with infected and non-infected customers. ISPs can also measure bandwidth utilization for different groups of customers.

Support costs can vary widely depending on how an ISP implements their program. Well-designed programs with clear information will have lower support costs than hastily built programs that are confusing to consumers.

Recommendation # 8: An ISP may provide remediation tools to the end user, either during or after the notification process. However, the ISP should not mandate that the end user run remediation tools. If the ISP provides tools to the end user, the end user should be allowed to exit the process without running any suggested tools or procedures.

Technology Considerations

ISPs that provide on-line tools to help the consumer to remediate their infection can be very effective and provide a high level of customer satisfaction. However, development of an automated, self-help system is non-trivial and requires testing for each family of malware infections. The working group is not aware of any ISP that currently develops its own software remediation tools—that task typically falls to anti-virus vendors and software developers.

When providing guidance and tools to the end user, the ISP must try to help the end user properly identify the infected device. Since most ISPs currently cannot see past the broadband modem in IPv4 configurations, the end user may not be on the suspected infected device when tools are provided. ISPs should provide as much information as possible to help the end user identify the infected device including such things as:

- Information on the operating system that the malware runs on. For example, many infections will only run on a specific operating system such as Windows, OSX, Linux, or Android platforms.

- Suggestions for improving security such as securing Wi-Fi access points to eliminate piggyback users unknown to the customer.
- Helping end users understand that new mobile devices could be the source of the malicious traffic. This would include visiting family and friends.

More sophisticated malware can block end users from accessing security vendors' websites and tools. Thus, ISPs may have to locally host third-party cleaning tools on their systems so consumers can access and run the tools.

Finally, it is recommended that the ISP review the cleaning procedure, possibly with the developer, to try to anticipate the end user's likely experience and identify potential problems. An example of potential problems includes malware that closes browser windows with specific security terms in their title so that consumers are unable to see the cleaning procedure. Another example is malware that disables a consumer's ability to access most security sites and their computer's security software. In this case, the consumer must run a locally-hosted version of a cleaning tool to disable the malware before complete cleaning and re-enablement of the computer security software is possible.

At a minimum, ISPs can host standard cleaning tools on their ISP website to assist subscribers with remediating their infection.

ISPs will need to integrate the systems or tools with their technical support websites and potentially with their network services to redirect customer traffic to appropriate content.

ISPs should review and integrate appropriate remediation steps for major malware families for infected subscribers with their online technical support information web pages. ISPs may elect to develop self-help remediation systems that walk a subscriber through the steps necessary to clean their infection.

If a self-help system is developed, the ISP will likely have to develop separate remediation pages for major families of malware. Different cleansing steps and tools will likely be required for different malware families. Close coordination with anti-virus and operating system vendors is important to ensure the tools adequately clean the subscribers' computers.

Currently, there are some startup companies who provide systems that will integrate into an ISP environment for subscriber infection remediation. Technologies vary depending on the architecture of the ISP's network, configuration of the subscriber's devices and networks, along with the level of integration into an ISP's network infrastructure.

Consumer/Market Considerations

Consumers who are successfully notified and remediate their infections indicate a high level of satisfaction in customer surveys. However, if the consumer has trouble identifying the infected device(s) or has trouble executing the remediation steps, their satisfaction drops quickly.

Human-factors testing of any remediation advice or system is critical for success. Most consumers do not have a strong technical background and information must be present in clear, non-technical terms and easy to follow instructions.

In addition, some more technical users will request additional information about how the ISP identified their systems were infected. The ISP should be prepared to handle these inquiries.

Infection remediation is the most intrusive of the five primary areas of Code participation. A well designed program can lead to greater customer satisfaction and low churn. A poorly designed program may lead to higher customer call volumes and lower customer satisfaction.

Because of the significant challenges and barriers associated with developing and maintaining a remediation program, many ISPs include a recommendation that customers seek professional help from computer retailers, computer support organizations, and the ISP's own professional services group for a fee. Although assisting customers by providing basic resources for a "do it yourself" approach is advisable, ISPs should make clear that remediation is ultimately the consumer's responsibility, and that they should seek qualified professional assistance if they are not qualified or capable of remediating their devices.

Once notified of an infection, most consumers have a strong desire to protect their personal information and clean the infected devices. It is only a minority of consumers that express no concern for the infection and wish to be left alone.

Operational Considerations

Many ISPs have a goal to minimize human customer support and provide clear, self-help procedures or third-party professional help for their customers. ISPs need to expect an increase in technical support calls for a minority of users.

Design and testing of remediation procedures may go a long way towards minimizing human support costs but actual testing of third-party remediation software may require significant resources from the ISP.

Depending on the design of the program, operational processes will need to account for customer notification and remediation campaigns. ISPs will need to consider whether they will provide multiple notifications for the same infection and what, if any, escalation procedures they will implement to encourage customers to fix their issues.

Procedures to redirect customers to third-party support services are the minimum level needed to help customers remediate their infections.

Well-designed programs integrate customer notification and remediation programs with their call center and chat support services. This improves the customer experience and allows customers to verify notifications sent to them.

Depending on the program, network engineers, information security engineers, system engineers, database administrators, corporate communications personnel, and human factor personnel may all be needed to implement a system.

Generally, new database systems are required to track notifications and remediation instructions. The system may need to integrate with the ISP authentication system and network elements to redirect customer traffic.

Lab testing, human factors testing, customer surveys, and infection reduction are some important measurements to ensure that an ISP program is working effectively.

Legal/Regulatory/ Policy Considerations

ISPs should consult with legal counsel regarding their remediation efforts. Care should be taken, for example, to disclaim responsibility for third-party software tools.

Financial Considerations

Once a system is in place, a small dedicated team may be required to support the program, depending on the size of the ISP broadband customer base. This may well be related to the overall size of the ISP implementing the program. A well designed system will allow new instructions and tools to be added as new botnets are identified and tracked.

System tuning and incremental upgrades should be supported, otherwise customer support costs may gradually rise instead of falling.

Some ISPs have documented cost saving primarily in support call volume reduction. Customer surveys indicate high satisfaction rates when clear notifications are provided and customers successfully remediate their infections.

ISPs can measure relative support costs from infected customers versus non-infected customers. Value added security services may also improve business economics.

Recommendation # 9: As part of the notification process, ISPs may wish to include guidance (depending on the nature of the bot in question) that settings on customer-owned network equipment such as home gateways and routers may have been altered and should be restored to a secure state, depending on the nature of the bot infection.

Technology Considerations

Although currently rare, infection of home gateways and routers can be very difficult to troubleshoot and identify. If a family of malware can alter these devices, end users should be given guidance on resetting these devices to the original factory settings, reconfiguring the devices and properly securing the devices (such as setting unique passwords on the device and upgrading the device to incorporate security fixes).

Ideally, existing ISP customer premises equipment (CPE) management systems would be used to mitigate this type of event. Otherwise, manual procedures are required to assist customers in remediation.

System integration with the bot notification and remediation system is probably not necessary.

The level of technical support varies greatly with the type of equipment infected. If ISP-supplied CPE equipment is infected, regular ISP broadband service could be compromised. Some ISPs have CPE management systems in place, minimizing the impact of infected device identification and remediation. If such systems are not in place, ISP remediation activities could have significant manual resource impacts.

Third-party network based equipment remediation can be very difficult for the ISP to troubleshoot.

Consumer/Market Considerations

For ISP supplied CPE, mitigation with little consumer interaction is best.

If ISP CPE is consumer owned, permission is probably needed to remediate the issue. This is important in case the remediation is not successful and impactful to the subscriber's broadband service.

Operational Considerations

ISPs may need to develop new operational procedures related to informing customers of ISP CPE configuration activity. For example, if a customer actively manages their CPE, an ISP reset of that device to factory settings without notifying the customer may cause that customer difficulty when he or she next seeks to access their device, or may erase customized security settings the customer has configured on their CPE.

Legal/Regulatory/ Policy Considerations

ISP activities in this area should be carefully reviewed with counsel to ensure compliance with applicable law, the ISP terms of service, and other similar policies.

Financial Considerations

Financial considerations for remediating CPE-based exploits are similar to the considerations discussed elsewhere for remediating other types of devices.

COLLABORATION

Recommendation # 10: ISPs, commercial security vendors, academia and government should collaborate to share

- **threat,**
- **attack,**
- **notification,**
- **education,**
- **measurement and**
- **mitigation**

methods and/or practices planned for or employed in ISP networks, and where practical an evaluation of their effectiveness.

Technology Considerations

Staff (e.g., security) familiar with multiple facets (e.g., detection, measurement, and mitigation) of the botnet issue need to participate actively in forums where issues relevant to botnet mitigation issues are discussed.

Industry organizations including but not limited to the Messaging, Malware, and Mobile Anti-Abuse Working Group (M³AAWG), the Internet Engineering Task Force (IETF), and On-Line Trust Alliance (OTA)) are typical of appropriate forums for collaboration activities.

Consumer/Market Considerations

Collaboration is not a consumer-facing activity. However, consumer-facing information may, as appropriate, indicate collaboration roles and/or practices.

Collaboration has been found to stimulate and motivate participants to higher levels of accomplishment.

It is generally accepted as fact that collaboration activities have had a key role in the management of a very similar problem – email spam. A large number of financially motivated organizations including service providers are now sponsoring and participating in similar anti-bot activities, indicating a near-consensus view that collaboration is an important component of botnet mitigation.

Operational Considerations

There are no mandatory impacts that would adversely impact operational processes. It is expected that collaboration would have a positive impact on the effectiveness of anti-bot operations.

Some automated systems that facilitate operational data exchange may be developed.

Legal/Regulatory/ Policy Considerations

Collaboration among competitors may require legal advice and guidance. Industry forums generally offer an appropriate forum and guidance for appropriate collaboration.

Financial Considerations

The benefits of a better-educated and empowered anti-bot ecosystem is empowering, reducing bots and bot mitigation costs.

The primary costs are for staff salaries, travel, and association memberships.

Appendix 4 – Metrics Guide



March 2013

Appendix 4

BOT and BOTNET METRICS Guide (Analysis & Recommendations)

U.S. Anti-Bot Code of Conduct (ABC) for Internet Services Providers (ISPs)

WORKING GROUP 7 - Botnet Remediation

Basis and Outline of This Guide

The CSRIC III *Working Group Descriptions and Leadership* document (last updated November 15th, 2012),⁵ provides among other things that "[Working Group 7] shall identify performance metrics to evaluate the effectiveness of the ISP Botnet Remediation Business Practices at curbing the spread of botnet infections."

Some issues examined in this Metrics Guide are currently out of scope for WG7, but are included since they are important issues that need to be addressed by the wider bot ecosystem and in future efforts.

This report is provided in fulfillment of that Metrics requirement, and has the following structure:

- I. Expected Audiences for This Metrics Guide
 - II. Thinking Precisely About What Is and Is not A "Bot"
 - III. What Sort of Botted "Things" Should We Be Trying to Count?
 - IV. Substantive Questions About Bots
 - V. Some Statistical Questions Associated With Botnet Measurements
 - VI. ISPs As A Potential Source of Botnet Data
 - VII. Sinkholing, DNS-Based Methods, Direct Data Collection and Simulations?
 - VIII. Recommendations
- Appendices

⁵ <http://transition.fcc.gov/pshs/advisory/csr3/wg-descriptions.pdf> at PDF page 7.

I. Expected Audiences for This Guide

While the primary audience for this bot metrics guide is the FCC CSRIC itself, it is not the *only* such audience. Many other audiences also need or want botnet metrics, including ISPs, the public, other federal agencies, Members of Congress, public interest organizations, members of the media, security software and security hardware vendors, law enforcement agencies, cybersecurity researchers, and governments abroad.

These various audiences may have widely varying bot metric information needs:

- (a) The CSRIC and the FCC may question, "Do we have empirical evidence that the Anti-Botnet Code is helping, or do we need to try some other approach? If it is helping, how much?" "Overall, can we say that the botnet problem is getting better, or getting worse?" "What anti-botnet strategies have proven to be the most (and the least) successful?" "When "improvements" are implemented in the Code, are they effective?" "Are US ISPs doing as well at tackling bots as their counterparts in Canada, Germany, Japan, Australia, etc.?"
- (b) ISPs considering adoption of the Code may wonder, "Is the Anti-Botnet Code for ISPs worthwhile? What might it cost us to participate? What benefits might accrue to us if we do so? Do the benefits exceed the costs?" ISPs that have already adopted the Code may wonder what techniques or products are most effective. How can we work towards agreed to metrics?
- (c) The public may want data that will allow them to choose an ISP that takes the bot problem seriously, and has taken effective steps to deal with bots targeting customers.
- (d) Other federal agencies may ask, "We have been fighting bots, too. How effective have the FCC's efforts been compared to ours? Are there opportunities for us to collaborate on targeted joint initiatives? If so, where is the low-hanging fruit?"
- (e) Members of Congress may need botnet metrics to determine if new legislation is required, or if existing legislation requires additional funding in order to be fully effective.
- (f) Public interest organizations may want botted users to be protected from botnet threats, but only in a way that is appropriate and privacy-respectful. A sense of the magnitude of the problem is critical to sizing up what might be necessary, and metrics also provide programmatic transparency.
- (g) Members of the media will be interested in understanding and reporting on government efforts and initiatives, and will want to see documentation about how much work is being done on botnets, and to what effect.
- (h) Security software and security hardware vendors may view botnet metric requirements as potentially driving new markets for new security gear -- or defining how well their existing gear works. In a competitive marketplace, metrics may define "winners" and "losers" and be important drivers for keeping existing customers and gaining new ones.

- (i) Law enforcement agencies may eagerly seek botnet metrics to help them to target and optimize their enforcement activities, endeavoring to target their limited cyber enforcement dollars in a way that gives taxpayers the most "bang for their buck."
- (j) Academic and commercial sector cybersecurity researchers might want access to raw empirical data about botnets for use in their own analyses and for identifying best practices in eliminating bots.
- (k) Governments overseas may look at our bot metrics to see if this program is something that they should be doing, too.
- (l) Other audiences in the bot ecosystem may have additional metrics needs.

An unfocused/*ad hoc* botnet metrics program is unlikely to serendipitously meet the requirements of all those diverse audiences. The focus of WG7 is to create metrics for Code participation. The metrics that are needed now -- and that may be needed in the future -- must be explicitly and carefully defined, or we run the risk of finding ourselves with no evidence with which to answer critical operational and policy questions relating to Code participation.

At the same time, we must remain cognizant of the fact that collecting and reporting data about bots is potentially burdensome, intrusive, and expensive. In addition, *ad hoc* data collection that is not done on a rigorous and consistent basis may lead to problems at the analytic stage, where data from multiple sources may turn out to be non-comparable, and thus unusable for comparative analysis. Therefore, any data that is targeted for collection should be data that is needed, and which will be collected consistently, and used in meaningful ways that justify the costs of its acquisition.

However, it should be noted that ISPs are at an early stage in terms of deployment of anti-bot technologies. Any guidance as to the collection of metrics should be offered in the light of the probability that mandating particular metrics may serve as a barrier and not incentive to rapidly deploy such platforms and capabilities.

A Specific NON-Audience: Botmasters and Other Cybercriminals: While there are many legitimate audiences that are welcome to industry botnet metrics, there is one explicit non-audience: botmasters (and other cyber criminals). We need to explicitly recognize that botnet metrics, done wrongly, have the ability to potentially help our enemies and undercut anti-botnet goals. A few simple examples:

- (a) Some ISPs may worry that if publicly identified as working diligently to combat bots, they may be targeted for serious and ongoing Distributed Denial of Service (DDoS) attacks by unhappy botmasters.
- (b) Giving detailed and accurate information about where and when bot activity was observed may be sufficient for a botmaster to identify (and subsequently avoid!) honeypots (or other data collection infrastructure) in the future. If that happens, valuable (sometimes irreplaceable) data collection sources and methods may be compromised.
- (c) If our botnet metrics include "per-bot" cleaning and removal statistics, botmasters might be able to use that feedback to learn what bots have proven hardest to remove, information that they can then use to "improve" future bots, making them harder to mitigate or remove. We do not want to teach our enemies how to better technically overcome our defensive strategies.

Our botnet metrics efforts must be structured so as to avoid accidentally benefiting our opponents or inhibiting our ultimate goals.

A Concrete Botnet Metrics Example From the Media: The following quote from the online cybersecurity industry news site Dark Reading⁶ nicely underscores why botnet metrics matter, and can be hard to do well:

Seven months after a coalition of government and industry organizations announced a broad set of voluntary guidelines [e.g., the ABCs for ISPs] to help Internet service providers clean their broadband networks of malware, the effort has yet to produce measureable results. [...]

So far there is no evidence that the effort is producing meaningful results. In the third quarter of 2012, for example, 6.5 percent of North American households had malicious software on at least one computer, according to the data from Kindsight's latest report. The rate is a slight increase from the 6 percent of households that showed signs of malware infection in the first quarter of the year.

That anecdote nicely illustrates some of the metrics-related challenges the industry faces:

- (a) All types of malware are treated as if they represented "bots," even though some of the most common types of malware are not even remotely "bot"-like. We need to be very precise about what is and is not a bot if we are to collect any sort of useful numbers.
- (b) By looking at population-wide (total) infection rates, the infection rates of Code-participating ISPs get comingled with the infection rates of non-subscribing ISPs. Given that comingling, an uptick in bots among non-participating ISP users might offset any improvement in the number of bots seen in participating ISPs' user populations.
- (c) That study looked at the infection rate for "North American"⁷ households, while the Code only targets U.S. ISPs and their residential customers.
- (d) That study looked at infection rates for *households*, rather than *devices*. There might be a half dozen devices in a household, but if even one is infected, the entire household will

⁶ "Anti-Botnet Efforts Still Nascent, But Groups Hopeful,"

<http://www.darkreading.com/security-monitoring/167901086/security/news/240143005/anti-botnet-efforts-still-nascent-but-groups-hopeful.html>

⁷There is often a tendency to treat "North America" as if it is just comprised of the United States and Canada (with everything else in the Western Hemisphere being part of Latin/South America and the Caribbean), but in fact there are actually 29 countries that are serviced by ARIN, the "North American" Internet number allocation organization. If a researcher determines what is a "North American" household by checking to see if the IP address associated with each infection came from ARIN (rather than some other entity, such as RIPE, APNIC, LACNIC or AFRINIC), the anti-botnet efforts of U.S. ISPs will be potentially conflated with the botnet experiences of 28 other countries or territories. That means that even if U.S. botnet numbers improved, domestic improvements (if any) may end up marginalized or eliminated by a hypothetical worsening of Canadian/Mexican/Caribbean/other "North American" countries bot numbers.

get flagged as bad. This can skew the proportion of a population that ends up getting reported as infected.⁸ This is a challenge.

- (e) On the other hand, what about infected devices other than just desktops or laptops? For example, what about smart phones and tablets? Are we also counting infections on those devices? What about other sorts of devices, such as "smart TVs" or Internet-connected gaming consoles?
- (f) Not all broadband customers (nor all infected broadband customers) are "households." For example, another important broadband customer segment might be small and medium-sized businesses and similar organizations (such as broadband-connected primary and secondary schools). The infections in those populations are "out of scope" for the WG7 charter.
- (g) What if a botted device is offline, and thus not "showing signs of an infection" (to use the language from the article). Does that/should that "infected but offline" device still "count"? Are we counting the infection or the effect of infection?
- (h) What constitutes a "meaningful" or "material" change for the better (or worse)? Is there some level that we may eventually reach (even if it is not zero) that we can all agree is "good enough?"
- (i) Improvements in bot detection can show increases in the bot population without any actual growth; apparent increases in botnet infection levels may actually just reflect improved botnet detection sensitivity.
- (j) There are multiple philosophies around infection notification. Some ISPs notify based on known bot malware with known cures, others notify to enable users to change behavior.

The answers to those and other questions largely shape the bot metrics space, and those choices largely determine the answer that one ultimately finds.

We need to address these and other issues if we are to be able to provide meaningful metrics about the state of bots in the United States, and if we are to be able to measure the potential impact of the Code. What are, or should be, the key performance indicators for Code participation?

Let us begin with the issue of what is or is not a bot.

⁸ To understand this distinction, imagine a hypothetical media report about the impact of the flu on 150 area businesses employing 30,000 people. If each business had exactly one employee sick with the flu (a total of 150 sick people among all area businesses), we could either report that *"100% of businesses had been hit by the flu"* (since each business does in fact have exactly one employee sick with the flu), or that *"just 1/2 of one percent of all employees have the flu"* (e.g., since $150/30,000 * 100 = 0.5\%$). These two different metrics convey radically different stories about the hypothetical flu problem in area businesses.

II. Thinking *Precisely* About What Is and Is not A Bot

What Exactly Is a Bot? In an earlier report,⁹ Working Group 7 provided a general definition of "what is a bot," stating:

A malicious (or potentially malicious) "bot" (derived from the word "robot" [...]) refers to a program that is installed on a system in order to enable that system to automatically (or semi-automatically) perform a task or set of tasks typically under the command and control of a remote administrator (often referred to as a "bot master" or "bot herder.") Computer systems and other end-user devices that have been "botted" are also often known as "zombies".

Malicious bots are normally installed surreptitiously, without the user's consent, or without the user's full understanding of what the user's system might do once the bot has been installed. Bots are often used to send unwanted electronic email ("spam"), to reconnoiter or attack other systems, to eavesdrop upon network traffic, or to host illegal content such as pirated software, child exploitation materials, etc.¹⁰

While that is a fine definition as far as it goes, it may not sufficiently emphasize one critically important point:

Not all malware is bot malware.

Characteristics that can be used to differentiate malware in general from bot malware in particular include:

Characteristic	Malware	Bot
1. The malicious software was installed without the user's informed consent (installation may have been done involuntarily, or the user may have been deceived into allowing installation of the malicious software)	YES	YES
2. The malicious software performs some sort of unacceptable or illegal activity (interferes with system use, compromises private information, sends spam, scans/attacks other systems, hosts illegal content, etc.)	YES	YES
3. The malicious software does <u>work over the network</u> after installation	POTENTIALLY	ALWAYS
4. The malicious software enables a remote administrator (the "botmaster") to <u>instruct the infected system to do specific work of the botmaster's choosing</u> (that is, the bot can be remotely "steered")	POTENTIALLY	ALWAYS
5. The unauthorized remote administrator is able to control <u>multiple</u> infected nodes as a <u>unified entity</u> , allocating malicious work across a set of worker nodes.	POTENTIALLY	ALWAYS

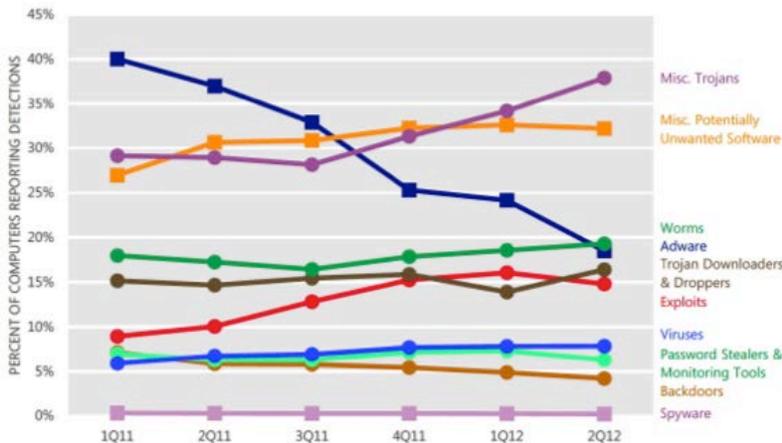
An illustrative/non-exhaustive list of current and historical malware families that are generally agreed to be "bot malware" can be found at <http://pages.uoregon.edu/joe/botlist/botlist.pdf>.

⁹http://www.maawg.org/system/files/20120322%20WG7%20Final%20Report%20for%20CSRI%20III_5.pdf

¹⁰ Bots are defined in IETF RFC6561: <http://tools.ietf.org/html/rfc6561>

An example of the challenges of determining what is and what is not bot malware can be seen in this graph from the Microsoft Security Intelligence Report (below),¹¹ which breaks out ten different types of malware. Bots are included mostly under “Backdoors”, but the report makes no specific mention of what is or is not a bot.

Figure 28. Detections by threat category, 1Q11–2Q12, by percentage of all computers reporting detections



Round markers indicate malware categories; square markers indicate potentially unwanted software categories.

A Botnet Malware Registry? To help eliminate ambiguity over what is and is not a bot, one option would be to create a voluntary botnet malware registry. An excellent foundation for a registry of this sort might be the site <http://botnets.fr/> which currently catalogs over 300 botnet families by name.¹² With the polymorphic nature of some bots, each device infection may be unique, but it still may be possible to identify the family the bot belongs to. Once an agreed upon bot registry is available, whether that is botnets.fr or something else, malware that has been found to be "bot" malware could then be listed in that registry. Such a registry would be voluntary and would likely have a only a partial view of active bots, nevertheless, it may be useful in sharing information about bot families.

While this might sound like a small step, it actually enables significant bot-related research. For instance, anti-malware vendors, when analyzing and cataloging malware they detect, could then potentially voluntarily add an "is this malware a bot?" attribute to their malware catalog entries (based on the registry), and potentially employ that attribute as part of their periodic malware reporting. For example, in addition to any other statistics an anti-malware vendor might share, an anti-malware vendor might also hypothetically report on:

- (a) The number of new bot malware families discovered that quarter,
- (b) The percent of devices seen infected with each of the dozen most significant bots, and
- (c) The total number of hosts detected as infected with one or more bots.

Having a common bot definition would allow multiple reports of that sort to be compared by all members of the ecosystem: do all anti-malware vendors see the same number of new bot malware families? Do they see approximately comparable new levels of infection? Until we agree on what is and is not a bot, it

¹¹ Microsoft Security Intelligence Report, Volume 13, <http://www.microsoft.com/security/sir/>

¹² While the primary language of that site is French, content is also available in English via the selector in the left hand column.

is impossible to tell if apparent differences are due to bot definitional differences, or other differences (such as a different customer base, differing detection efficacy, etc.)

"If It Acts Like a Bot:" In some cases (for example, in the routine case of an ISP that does not have direct administrative access to a customer's system), if a system exhibits empirically observable "bot-like behaviors" (such as checking in with a botnet command and control host,¹³ or spewing spam, or contributing to a DDoS¹⁴ attack), even if a particular bot cannot be identified, the system might still get tagged as being botted.

Tagging botted systems based on their externally observable behavior may be necessary when direct access to the device is not possible, but also in cases where devices are infected with malware that is so new that antivirus companies have not yet had time to identify that malware.

Therefore:

If a device acts like it is infected by a bot, even if it cannot be identified as infected by a particular type of bot malware, it may be necessary to assume it is botted.

III. What Sort of Botted "Things" Should We Be Trying to Count?

If consensus is reached on what is and is not a bot, we are a large part of the way to being able to ask meaningful/ measurable questions about them.

However, we also need to decide one other critical issue, and that is deciding precisely what sort of botted "things" we are going to count.

What I am Able to Measure Will Depend on My Role In the Ecosystem:

- (a) If someone were to go "boots on the ground" and actually check all the devices in a number of households to see whether any device is infected, those researchers would have the option to count individual infections, or botted systems, or botted IP addresses, or botted users, or botted households. While going "boots on the ground" might seem to provide the most flexibility and most comprehensive data collection options, it is also the most potentially expensive option, and it presumes access to a household's devices, access which might be viewed as intrusive and routinely denied.
- (b) On the other hand, if an ISP detects bots based on malicious network activity associated with a particular IP address, the ISP is likely going to count botted IP addresses or botted subscribers.
- (c) If an antivirus company or an operating system vendor scans/cleans end-user devices, it is going to end up counting individual infections,¹⁵ or perhaps botted devices.¹⁶

¹³ A "command and control host" is a system that a botmaster uses to run his botnet.

¹⁴ A DDoS attack is a Distributed Denial of Service attack, often conducted by flooding a site with so much bogus traffic that the site's network connection or servers cannot keep up, thereby preventing legitimate users from being able to use that site.

¹⁵ A single system might have multiple simultaneous infections.

¹⁶ One user might have two or three systems, or one system might be shared by multiple users.

- (d) If a survey research company calls people up on the phone and ask, "Have you ever been infected with 'bot'?" those survey researchers are likely going to end up counting botted users.¹⁷

Different parties will contribute different views of the problem. While those views may be different, all are potentially valuable and important.

Desktops and Laptops Only? Are we going to measure all kinds of botted devices, or are we just going to count botted desktops and laptops? For example, consider smart phones in particular. The number of smart phones is now material, and malware is increasingly attacking and infecting at least some types of those devices.¹⁸ Users also have a growing number of tablets, Internet-connected "smart TVs" and set-top boxes, gaming consoles, and other devices that may be targeted for compromise. Should all those sorts of devices be counted, if botted? We think so, yes. To ensure a comprehensive bot "picture," we suggest that any program of bot metrics should include all types of Internet-connected devices, but, as that data is collected, it should include the type of device involved, thereby allowing analysts the option of reporting about all devices, or just some particular subset of all devices, such as just traditional laptops and desktops, or just smart phones.

Online Devices Only? We must recognize that in most cases we can only count botted devices that are "live" or "that we can see" on the network. Other devices may be botted, but remain undetected/unknown as a result of how we detect botted hosts.

An easy way to understand this is to imagine that we are an enterprise that actively scans its corporate systems with Nessus¹⁹ or a similar security scanning tool in an effort to identify devices that appear to be botted. Obviously, if a device is not connected to the network, or is not powered up when our network scan takes place, a potentially infected device will not be able to be found with that tool.²⁰

Similarly, if an ISP identifies botted hosts based on spam (or other network artifacts visible to the ISP "on the wire"), a botted host that is offline (or in a walled garden) will not be seen "making noise" or "causing problems" on the ISP's backbone²¹ where instrumentation exists, and also will not/cannot be noted as being botted by external parties relying on externally visible symptoms to tag devices as being botted.²²

¹⁷ If we talk to three different people from the same household, and they all used the same botted computer, we might potentially get three reports that are all about that single botted system. This scenario also runs afoul of multiple other issues, including things as basic as the fact that users may not know what a bot is, or they may forget having been botted.

¹⁸ "Virtually All New Mobile Malware is Aimed at Android,"
<http://www.androidauthority.com/mobile-malware-aimed-android-112403/>

¹⁹ http://en.wikipedia.org/wiki/Nessus_%28software%29

²⁰ This raises an interesting methodological question: if we scan and cannot reach a potentially botted host, how often should we attempt to rescan it? Once? Twice? Time after time after time? Never?

²¹ A related potentially important methodological question (at least from the point of view of ISPs actively mitigating botted hosts): if a botted system has been detected as being botted and successfully put into a so-called "walled garden" where it cannot cause problems for other Internet users, should that host still be counted as "botted"? Or do we need an additional category to capture systems in this status, "botted but not online and not able to cause problems," perhaps?

²² This may be a material problem, kin to giving cough syrup to lung cancer patients. You may stop the externally visible symptoms with symptomatic treatment, but in doing so, you are not

Measurement Window-Related Decisions Will Have A Material Impact on Bot Detection Rates: If a botted device is only online occasionally, *when* and *how long* we look for bots (e.g., our "measurement window") may strongly impact how many bots we find.

For instance, if an ISP is detecting bots based on characteristic botnet network activity, when/how long the ISP collects bot-related network flow records will strongly influence botnet detection rates. To understand why, note that if we only watch for botted hosts during a brief window during the business day, we might miss home systems that are turned off except when a family member is using them at night. On the other hand, if we collect bot data during evening "prime time" hours, we will likely miss any botted work systems that may only be on and in use during the normal 8 to 5 work day. Therefore, if we try to count botted hosts during too brief a time period, we may miss some bots.

If we go to the other extreme, and *continually* watch for botted hosts, we will virtually certainly see the same botted host more than once, and since we may have trouble telling one botted host apart from another, we run the risk of counting a single bot more than once, simply because in most cases there is no unique identifier that we can use to track a particular botted host from one sighting to the next.

Unique Identifiers: If each botted host *did* have a unique identifier, we could collect data over a protracted period and not have to worry about counting the same botted host multiple times. Unique identifiers for botted hosts would also greatly simplify the process of aggregating (or "rolling up") fine grained records appropriately. For example, we could tag a record about each individual infection with the unique identifier associated with that botted host, and then we could easily consolidate that data if/when we wanted to do so. We note that despite these potential benefits, uniquely identifying machines on the internet raises a host of privacy concerns and is currently an exceedingly difficult and resource-intensive goal to attempt to achieve.

Unfortunately, if we do not or cannot use unique identifiers, our measurements may end up profoundly flawed.

Consider one example mentioned in the ENISA report on *Botnets: Detection, Measurement, Disinfection & Defence*:²³

curing the underlying disease. Sometimes having bad symptoms can be a good thing when it comes to forcing attention to be paid to a serious underlying problem.

²³ *Botnets: Detection, Measurement, Disinfection & Defence*, European Network and Information Security Agency, 7 Mar 2011, <http://www.enisa.europa.eu/activities/Resilience-and-CIIP/critical-applications/botnets/botnets-measurement-detection-disinfection-and-defence>

In a report about the temporary takeover of the Torpig botnet [50], different measurement methods could be applied for 10 days and compared:

Sinkholing of botnet traffic provides a view of the botnet's live population.

- 1,247,642 unique IP addresses.
- The number of observed IP addresses increased almost linearly over time.
- 182,800 hosts estimated by unique identifiers.
- 75% of unique identifiers were detected during the first 2 days.

Taking the unique identifiers as a reference, the IP addresses yielded an overestimate factor of 6.8. This, and the fact that the bot's IP addresses varied over the time, leading to a constant increase in observed addresses, are an indicator of the unreliability of measuring IP addresses only.

Trusting size estimates based on such identifiers can also be open to question questionable, in case the generation algorithm is flawed or has been manipulated by the botmaster. For example, in IRC botnets, bots use nicknames as identifiers. Depending on the botnet, these nicknames are often generated every time the computer is

Taking observed unique IP addresses as the only indicator of a botnet's size usually leads to drastic inaccuracy and often to overestimations.

Another Reason Why IP Addresses Are Not Unique -- Network Address Translation (NAT): Another potential complication when it comes to mapping bots to IP addresses is the use of NAT. NAT is a public IP address-conserving technology that allows multiple devices to share a single public IP address. Because multiple devices share a single public IP address, malicious traffic from multiple devices may appear to come from just one IP address, resulting in an underestimate of the number of truly infected devices.

IPv6 Addresses: As the Internet runs out of traditional IPv4 network addresses, ISPs are beginning to use IPv6 addresses to supplement rapidly depleting stocks of IPv4 addresses. IPv6 significantly complicates the process of measuring botnets via their network traffic. Let us just mention a few of many reasons why this is true:

- (a) Many ISPs may not have IPv6 network flow monitoring that is on par with their IPv4 monitoring capabilities. These IPv6 network flow monitoring deficiencies may leave many ISPs fully (or at least partially) "blind" when it comes to monitoring IPv6 network activity, including IPv6 botnet-related network activity.
- (b) In other cases (as when ISP customers obtain tunneled IPv6 connectivity from a third-party provider), the customer's "home" ISP may have little or no visibility into the content of IPv6 tunneled traffic. Any third-party network researchers taking network measurements would see that customer's IPv6 traffic (including any IPv6 *bot* traffic) emerge from the tunnel-provider's network infrastructure, not the home-ISP's network infrastructure. That would obviously complicate any global botnet measurement work involving IPv6 connectivity.
- (c) IPv6 network address assignment technologies also play a role in potentially simplifying -- but more likely complicating -- botnet measurement efforts. To understand how this can be true, understand that in IPv6, IPv6 addresses may be assigned in multiple different ways, including, but not limited to ...

- (i) *Static IPv6 address*: If a botnet host uses a static IPv6 address, it is easy to track it over time, just as it is easy to track a botnet host that is using a static IPv4 address. Static IPv6 addresses are expected to be rare, however, except for manually configured servers.
- (ii) *Stateless Address Autoconfiguration (SLAAC)*: Hosts may have an automatically assigned²⁴ IPv6 address that leverages a reformatted version of the system's hardware MAC address. When devices use addresses of that sort, it becomes possible to potentially track a botnet host over time and even across multiple ISPs. This is an example of how IPv6 addresses can simplify botnet measurement activities.
- (iii) *Privacy Addresses*: IPv6 hosts can also use constantly changing IPv6 "privacy addresses."²⁵ Hosts using IPv6 privacy addresses periodically and automatically change their IPv6 addresses in an effort to make it harder for those systems to be systematically tracked by marketers or others. If a botnet host is using IPv6 privacy addressing, it becomes very difficult for researchers to accurately follow the IPv6 addresses that that system may be using over time, and as a result, we may potentially dramatically over-estimate the actual number of IPv6-connected bots.
- (iv) *DHCPv6*: In a fourth scenario, IPv6 using customers may receive IPv6 addresses via DHCPv6, the IPv6 analog of DHCP for IPv4. When DHCPv6 is used, ISPs can map a particular IPv6 address to a particular customer, but doing so remains somewhat awkward, and will not typically be practicable for large numbers of botnet hosts.

Counting Infections vs. Counting Botnet Devices: Let us set aside IPv6 issues for now, and just contrast two different devices:

Device A: Infected with one bot.

Device B: Infected with seven different bots.

Should each of those devices be counted as "one" infected device? Or should a botnet analyst generate one measurement for Device A, and seven measurements, one for each bot infection on Device B, thereby allowing each individual type of bot infection to be separately tracked?

This may be particularly relevant on large shared systems, such as high density web hosting sites, or "timesharing" Unix boxes with thousands of shell accounts, all sharing a single IP address. In that case, a single system/device might have multiple independent users, and multiple independent bot infections (including potentially multiple copies of the same bot!), all running in parallel.

Not All Botnet Hosts Are Equally "Potent": To understand what is meant by this, consider two hypothetical botnet hosts:

²⁴ IPv6 Stateless Address Autoconfiguration, <http://tools.ietf.org/html/rfc4862>

²⁵ Privacy Extensions for Stateless Address Autoconfiguration in IPv6, <http://tools.ietf.org/html/rfc4941>

Device C: An ancient consumer system connected via a legacy 56Kbps dialup connection

Device D: A high end server with multiple CPUs/multiple cores, lots of RAM, and gigabit Ethernet connectivity

Should each of those devices simply be counted as one "botted host"? There would certainly be a huge difference when it comes to the amount of spam or the volume of DDoS traffic that each of those two devices might respectively deliver. Does that mean that we should weighting botnet detections by some measure of capacity, such as their average spam throughput, or their average DDoS output?

Some Apparent "Bot" "Hits" May Not Be Real: For example, imagine researchers investigating a botnet: it is conceivable that they might attempt to "register" or "check in" fake "bots" of their own creation in an effort to figure out how a botnet operates, sometimes in substantial volume. In other cases, imagine an antibot organization that is attempting to proactively interfere with a bot by "poisoning" it with intentionally bogus information about fake bottled hosts. If we are measuring bots by counting the hosts that "check in" rather than the bots that are actually seen "doing bad stuff" we run the risk of overestimating the number of "real" bots that actually exist.

The M³AAWG program for bot metrics: The recently established M³AAWG malware metrics pilot program focuses on "subscribers" as the unit of analysis.

Because the number of subscribers will fluctuate over the course of a typical month, M³AAWG decided to use the number of subscribers as of the last day of the month.

ISPs participating in the M³AAWG data collection report the number of unique subscribers that have been found to be infected one or more times during the month. (What is/is not an infection is not explicitly defined, except to say that it should be an "infection" that is serious enough to motivate the ISP to contact the user about it.)

Participant ISPs will also report the number of unique subscribers that have been notified of a problem by whatever means (Short Message Service (SMS), phone call, email, web redirection/browser notification etc.) Multiple notices to the same subscriber count as one. That does not imply that the subscriber has read/received the notice. Subscribers often receive multiple communications from their ISP during the year and as such, it can be a challenge to get attention paid to a bot warning.

Given those values, one can compute the percentage of subscribers that have been found to be infected one or more times during a given month, and the percentage that have been notified of that infection.

As metrics, these implicitly involve accepting some compromises, e.g.:

- (a) Given the definition of these metrics, we cannot talk about how many infected customer devices may be present, nor how many distinct infections were seen, nor can we talk about whether a particular customer was repeatedly re-infected, or if the infection was on a laptop, smart phone, gaming console, etc.
- (b) The M³AAWG program does not focus solely on bots, since many ISPs want to protect their customers from all types of serious malware infections.
- (c) Choice of a month-long window means that day-to-day or week-to-week infection trends will not be able to be identified.
- (d) There are many other potential measurements related to infected customers that are not getting reported (for example, how much customer effort was required to disinfect and harden a typical infected system?).

These and other limitations were explicitly recognized and accepted by M³AAWG as part of its pragmatic program design decisions, recognizing that if it made the malware metrics reporting program too difficult or too time-consuming or too complex, many ISPs might be deterred from participating. Keeping the program simple and easy to participate in, may increase the number of ISPs willing to participate.

Another example of a pragmatic measurement choice was the decision to focus on the number of unique customer detections and customer notifications rather than the number of customer devices that have been cleaned up or rebuilt. (Because customers may use third-party services to clean up or rebuild their devices, ISPs may not know if a customer's device has been cleaned up, rebuilt, replaced, or remains infected (but is offline/dormant)). Establishing that remediation has occurred successfully is a major challenge. The manufacturers of remediation tools are generally unwilling to share data on effectiveness. Often privacy is cited as a reason. Additionally for the majority of remediation tools, no feedback loop exists from the tool back to the ISP to allow for measurement of remediation events.

IV. Some Substantive Questions About Bots

The M³AAWG program will yield some botnet related metrics. However, what are the other substantive questions about bots that we might like to answer?

What is the Order of Magnitude of the Bot Problem? If botted hosts are rare, we likely do not need to worry about them. On the other hand, if ISPs are being overrun with botted hosts, we ignore all those botted hosts at our peril.

If we do not (or cannot!) at least roughly measure botnets, we will not know if bots are a minor issue or a huge problem, and if we do not know roughly the size of the problem, it will be impossible for industry or others to craft an appropriate response.

Note that when we talk about "order of magnitude," we are not talking about a precise measurement. Rather, we are just asking, "Are 10% of all consumer hosts botted? 1% of all hosts? 1/10th of 1% of all hosts?" etc.

One example of such an estimate can be seen in Gunter Ollmann's "Household Botnet Infections:,"²⁶

Out of the aggregated 125 million subscriber IP addresses that Damballa CSP product monitors from within our ISP customer-base from around the world, the vast majority of those subscriber IP's would be classed as "residential" — so it would be reasonable to say that roughly 1-in-5 households contain botnet infected devices. [...] Given that the average number of devices within a residential subscriber network is going to be greater than one (Let us say "two" for now — until someone has a more accurate number), I believe that it is reasonable to suggest that around 10% of home computers are infected with botnet crime ware.

There are 81.6 million US households with broadband connectivity as of 10/2010.²⁷ If 20% of 81.6 million US broadband households were actually to be botted, that would imply that there are 16 million+ bots in the US alone. On the surface, this number appears high.

Let us consider another estimate, from the Composite Block List ("CBL"). On Sunday December 9th, 2012, the Composite Block List knew about 174,391 botted host IPs in the United States.²⁸

There are 245,000,000 Internet users in the US as of 2009 according to the CIA World Fact Book.

$174,391/245,000,000 * 100 = 0.0711\%$ of all US Internet users are potentially botted, [assuming 1 device/user].

Worldwide, that puts the US near the bottom of all countries, in 149th place on the CBL. On a per capita-normalized basis, that means that the US is among the least botted of all countries as measured by the CBL.²⁹

²⁶ http://www.circleid.com/posts/20120326_household_botnet_infections/

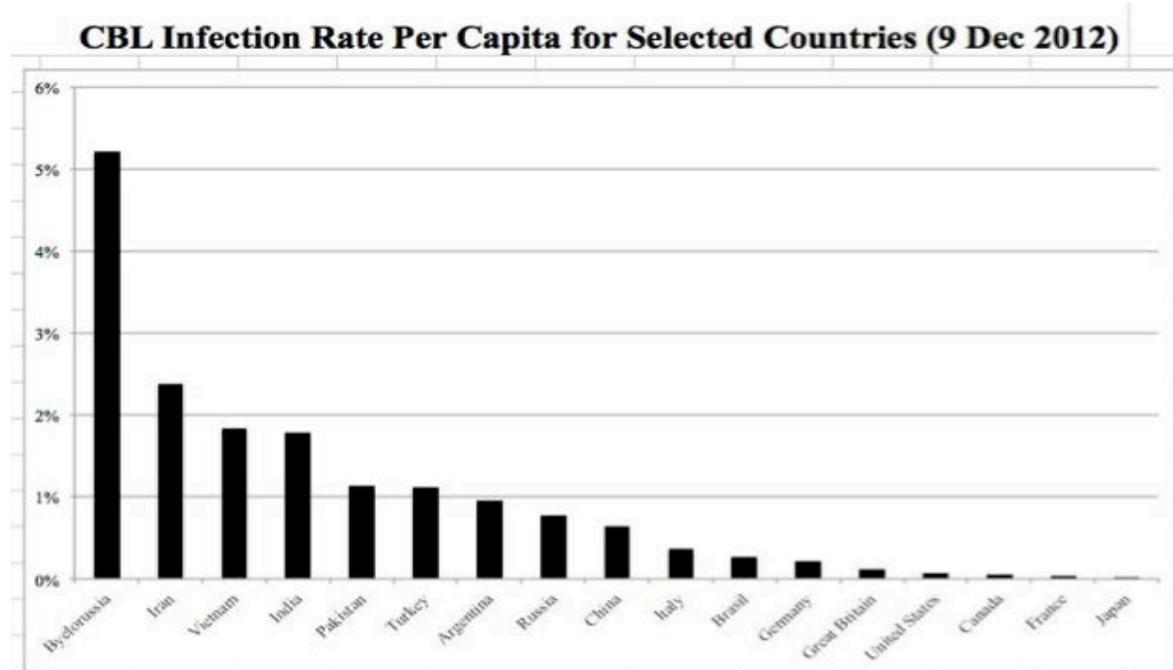
²⁷

http://www.census.gov/compendia/statab/cats/information_communications/internet_publishing_and_broadca

[sting_and_internet_usage.html](http://www.census.gov/compendia/statab/cats/information_communications/internet_publishing_and_broadca) (table 1155)

²⁸ <http://cbl.abuseat.org/countrypercapita.html>

See the following graph:



In fact, if there are only 175,000 bots here in the U.S., botted hosts have effectively become a "rare disease" in that when it comes to traditional medicine, the U.S. definition for a rare disease is one which afflicts fewer than 200,000 people in the United States.³⁰

How Could Those Two Estimates Be So Vastly Different? We think that there are two main reasons for the discrepancy:

- (a) The two estimates count different sorts of things (households that are detected as being botted vs. botted IPs seen sending spam)
- (b) The two estimates measure different populations: users worldwide who are connected via ISPs with enough of a bot problem that those ISPs have been motivated to purchase a commercial network security solution vs. users in the United States (where a real push to control bots has been underway for years).

Another Very Basic Question: How Many Different Families of Bots Are Out There? That is, are there three main types of bots actively deployed right now? Thirty? Three hundred? Three thousand? The proposed malware registry should help us to answer this question...

²⁹ Wonder which nations are the worst? At the time this report was drafted, and looking just at countries with 100,000 or more listings, the most-botted countries are Byelorussia (137,658 listings, with 5.2% of its users botted), Iraq (196,046; 2.4%), Vietnam (431,642; 1.85%), and India (1,093,289; 1.8%).

³⁰ <http://rarediseases.info.nih.gov/RareDiseaseList.aspx>

The number of unique types of bots is important because it tells us much about how hard it might be to get the "bot problem" under control. If there are only a handful of major bots, concerted effort should allow government authorities to shut them down, if the government makes doing so a priority. Conversely, if there are three thousand different types of bots out there, getting all those bots under control would be far harder.

Closely related to the question of how many types of bots are out there, how many botmasters are out there?

We might expect that the number of botmasters would roughly track the number of unique bots, but one bot "code base" might be "franchised" and in use by multiple botmasters, or one botmaster might run multiple different bots, eroding a direct one-to-one relationship between the two metrics.

In addition, we need to be mindful that the "franchising" of bot code means that one code base could be used multiple times, creating entirely separate botnets. Consider that one "bot," for instance Zeus, can have many different versions and there can be many Zeus botnets in operation simultaneously by different threat actors. Though this is only one "type" of bot, the actual scope of the problem it creates could be a great deal more significant than simple typologies would accurately reflect

Are There Any Trends Relating to Bots? That is, in general, is the bot problem getting better or worse over time?

Some anti-malware companies are already sharing data of this sort, at least for some types of bots. See for example the following graph from McAfee for the United States:³¹



Do Bots Show Any Sort of Operational Patterns? For example, hypothetically, does most botnet spam get sent "overnight" when US anti-spam folks are asleep but Europeans have already woken up? Does the number of bots increase during the weekend, and then go back down during the week? (This might be the case if a regularly employed botmaster just ran his or her botnet as a way to supplement his or her income on weekends, or if fewer anti-botnet people were paying attention/whacking bots on weekends.) Does the number of bots increase at the start of the month when people get paid and have money to buy spamvertised products, or does it peak in the month before Christmas (when people are most likely to be

³¹ McAfee Threat Report, Third Quarter 2012, PDF page 30,
<http://www.mcafee.com/ca/resources/reports/rp-quarterly-threat-q3-2012.pdf>

Christmas shopping), perhaps? If law enforcement arrests a botmaster or takes down a botnet, can we see a noticeable drop in the amount of spam sent, or do other botnets immediately step up and fill that now-vacant niche in the bot ecosystem? Here is one interesting graph of that last sort...:³²

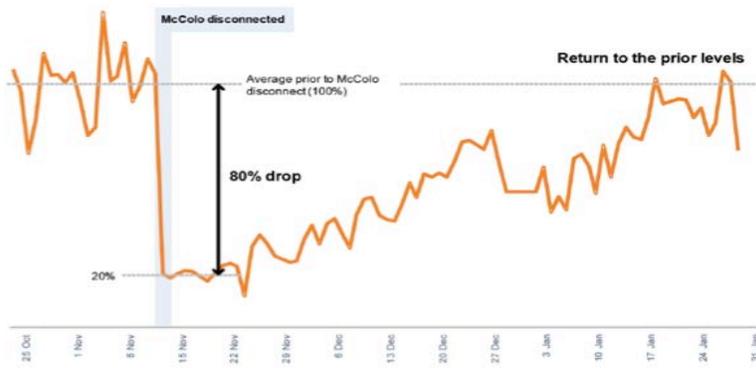
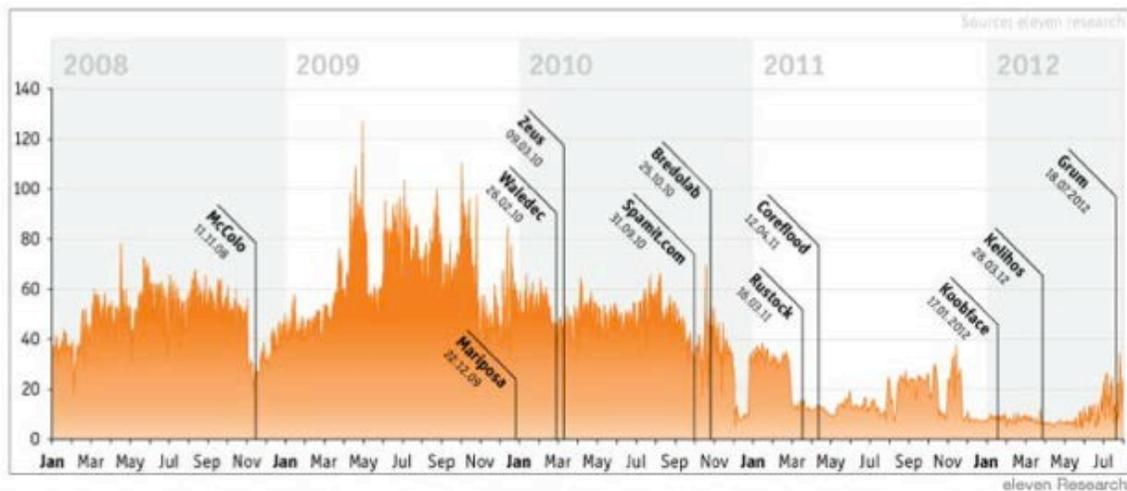


Figure 13: Temporary impact of the shutdown of hosting provider McColo on spam e-mail (Diagram by MessageLabs, Symantec Hosted Services). [2009]

Another example of an interesting long term botnet graph:³³

Botnets and Spam Development

Shutting Down Botnets and the Effect on Spam Volume



Spam volume and botnet shutdowns

³² *Botnets: Detection, Measurement, Disinfection & Defence*, European Network and Information Security Agency, 7 Mar 2011, <http://www.enisa.europa.eu/activities/Resilience-and-CIIP/>

critical-applications/botnets/botnets-measurement-detection-disinfection-and-defence

³³ <http://www.eleven.de/botnet-timeline-en.html>

HOW Are Botnets Being Used? While bots can be rapidly reconfigured from one purpose to another, do we have a clear understanding of how bots are currently being used, or how they have been used over time? That is, what fraction of all botnet capacity is used for: ...

- (a) Sending spam? (And how many spam emails are emitted from each individual botted host? Are those bots running "flat out," or just "loafing along?")
- (b) Participating in DDoS attacks?
- (c) Scanning network-connected hosts for remotely exploitable vulnerabilities?
- (d) Hosting illegal files?
- (e) Stealing private information?
- (f) Cracking passwords, mining BitCoins, or performing other compute-intensive tasks?
- (g) Are there bots that are installed but totally idle? If so, why? Excess capacity?

Understanding how bots are being used will help us to figure out we how we should try to measure bots.

For example, if bots are not longer being widely used to send email spam, we should not attempt to measure botnet populations based on the amount of email spam we observe. We should note that these kinds of judgments will be very dependent on the visibility level on the threat. Spam and DDoS bots, because of how they get used, tend to be noisy and thus are easy to spot. Credential harvesting bots and espionage bots, on the other hand, are intentionally much stealthier and harder to detect, and thus may be underreported in some summary bot typologies.

How Bots Are Being Used May Change Who Is Interested in Them: Hypothetically, if bots are no longer in widespread use for spamming, anti-spammers may lose interest in bots. On the other hand, if bots start to be widely used to conduct distributed denial of service attacks against critical government sites or critical national infrastructure, that change might increase interest in bots in the homeland security and national security communities.

We really need to understand/monitor the botnet workload profile as seen "in the wild," recognizing that this can change as quickly as the weather.

"Comparatively Speaking..." Another set of potentially interesting botnet metrics are comparative metrics:

- (a) Are American devices getting botted more (or less) than Canadian devices, or devices in Great Britain, France, Germany, Japan, Russia, China, Brazil, India or _____?
- (b) Not all countries are the same size. Should we normalize botnet infection rates by the population of each country (or by the number of people in each country who have broadband connectivity?)
- (c) Are all ISPs within the United States equally effective at fighting bots, or are some doing better than others? For example, if an ISP adopts the voluntary Code do they have fewer bots than other ISPs that do not adopt it?

- (d) Are there other important comparative differences that we can identify? For example, are older users (or younger users) more likely to get botted? Does it seem to matter what antivirus product or web browser or email client or operating system people use?

Comparative Raw Bot Levels Per Country From the CBL:³⁴

Country	Listings	%total	% Total Listings	%cumulative Total Listings	Rank
Total	8484086	100			2
CN	2512466	29.61	29.61	29.61	1
IN	1093289	12.89	12.89	42.50	2
VN	431642	5.09	5.09	47.59	3
RU	317034	3.74	3.74	51.32	4
TR	304213	3.59	3.59	54.91	5
PK	230561	2.72	2.72	57.63	6
BR	206180	2.43	2.43	60.06	7
IR	196046	2.31	2.31	62.37	8
US	174391	2.06	2.06	64.42	9
DE	143483	1.69	1.69	66.12	10
BY	137658	1.62	1.62	67.74	11
AR	130595	1.54	1.54	69.28	12

China looks pretty bad in that list, but then again, remember, China is a big country. How do they look, comparatively, once we adjust for their population?

Selected CBL Listings By Country, Normalized Per Capita:

CBL Infection Rate Per Capita for Selected Countries (9 Dec 2012)

Byelorussia	5.20840%			
Iran	2.38673%			
Vietnam	1.84604%			
India	1.78240%			
Pakistan	1.12849%			
Turkey	1.11707%			
Argentina	0.95367%			
Russia	0.77604%			
China	0.64588%			
Italy	0.36149%			
Brasil	0.27135%			
Germany	0.22032%			
Great Britain	0.11528%			
United States	0.07118%			
Canada	0.05990%			
France	0.04129%			
Japan	0.01595%			

³⁴ <http://cbl.abuseat.org/country.html> as of Sunday, December 9th, 2012.

Once we have normalized per capita, China is no longer leading the list (That dubious "honor" now goes to Byelorussia), however, China is still fully an order of magnitude more botted than the United States, even if China is fully an order of magnitude less botted than Byelorussia.

Pre/Post Longitudinal Studies: Given that it may be difficult to compare bot-related statistics collected by ISP A with bot-related statistics collected by ISP B, another option might be to track botnet stats longitudinally, within an individual ISP, over time.

For example, assume one would like to know if an ISP has fewer botted customers after adopting the Code than before (this is what some might call a "pre/post" study). One could then try to collect pre/post bot metrics.

Drilling Down on a Per-Bot-Family Basis: In addition to measurements made about overall bot infectivity, we also need the ability to "drill down" and get more precise estimates on a bot-family-by-bot-family basis, ideally both currently and historically. Per-bot-family measurements might include the number of devices infected with each particular major bot, and also related measurements such as:

- (a) the amount of spam attributed to each particular spam botnet
- (b) the volume of DDoS traffic attributed to each DDoS botnet
- (c) the number of command and control hosts that a bot uses
- (d) the geospatial distribution of hosts infected with each bot

Micro as Well as Macroscopic Measurements: Not all metrics are macroscopic measurements related to botnet infection rates. Some measurements of interest might be per-system micro values:

- (a) What does it cost to rent a bot on the open market?
- (b) How long does it take, and what does it cost to de-bot a single botted host? What factors make a device take more or less time to de-bot? Can we build a standardized cost model? For example, what is it worth to have a clean backup of a botted device? Does that make it significantly easier to get a botted device cleaned up and hardened?
- (c) When a device is found to be botted, does it tend to be botted with just one type of bot? If co-infections are routinely found, can we identify "clusters" of bot malware that are routinely found together, so that an anti-malware technician can then be told, "If you find bot A on a system, also be on the lookout for bot B, too?"
- (d) If a user is botted once, does that make them more (or less) likely to get botted again? That is, can we expect that that a once-botted user will become less likely to be re-botted as a result of that presumably unpleasant experience? Or are some types of users just inherently more prone to get reinfections, perhaps because of a failure to apply available patches, or inherently risky online activity patterns? If users do become reinfected, how long does it take?

Looking at This From A Different Direction: How Long Will A Typical Bot Live? Hypothetically assume that we are running a blacklist, and we list the IP addresses of botted systems when we see those devices send spam or check in with a Command and Control server that we are monitoring.

If you do not observe any subsequent activity from a botnetted and blacklisted device, when could we "safely" remove it? After a day? After a week? After a month? After 90 days? Never?

Some botnet blocklists deal with this issue by simply rolling off the oldest entries after the list reaches some target maximum size. (After all, if the device turns up again, it can always be relisted.)

Measuring Botnet Backend Infrastructure: While we have been talking about botnetted end-user hosts, another potential target for measurement is botnet backend infrastructure, such as botnet command and control hosts.³⁵ Potentially one could also track authoritative name servers associated with bot-related domains, and sites known to be dropping bot malware, and a host of other botnet-related things (beyond just focusing on botnetted hosts themselves).

A philosophical aside: Is there any risk that focusing on backend botnet infrastructure (including potentially doing Command & Control takedowns) will result in interference with ongoing legal investigations or metrics collection? If third parties do not target botnet backend infrastructure, can the Internet community be confident that law enforcement will in fact track and take down those botnet-critical resources? Are there ways that we can de-conflict this work without compromising operational security? This is a serious risk and must be dealt with appropriately. In many instances law enforcement may be actively pursuing a means to neutralize the threat that does not entail actually shutting down or seizing back-end infrastructure. We recommend establishing contacts for de-confliction as soon as possible.

V. Some Statistical Questions Associated With Botnet Measurements

How Precise Do Our Answers to These Questions Need To Be? "High precision" answers cost more than "rough" answers. (Think of this as the width of a confidence interval around a point estimate). If one wants to estimate a value within +/- 10%, it requires less work than if one wants to know that same value within +/- 5% or even +/- 1%. Exactly how precise do our measurements need to be, and why?

How Much Confidence Do We Need That Our Estimates Include the Real Value? For example, if we need 99% confidence that our estimate includes the real value for a parameter of interest, we can achieve that, but it might require accepting broader bounds around an estimate (or drawing more observations) than we would need if we could live with just a 90% level of confidence.

Notice the interaction between (a) the required precision, (b) the required confidence, and (c) the cost of obtaining those answers (typically the number of observations required).

Most people want high precision and high confidence and low cost, but you cannot have all three at the same time.

Budget: We really need to emphasize that if bespoke hard numerical answers to questions about botnets are needed, it is going to cost money to obtain those values. How much are we willing to spend to get those answers? Additionally, who is going to pay for this data to be collected?

If the answer is "zero," then we would suggest that in fact all our substantive questions about bots are just a matter of simple curiosity, and not something that is actually valuable ("value" implies a willingness to pay).

³⁵ See for example Zeus Tracker, <https://zeustracker.abuse.ch/>

If we also do not have a budget for data collection, our ability to rationally set the required level of precision (and the required confidence in our estimates) will also be impaired.

VI. ISPs As A Potential Source of Botnet Data

The CSRIC WG7 metrics presumption has inherently been that ISPs themselves might be a potential source of bot data about their botnet customers. While this is an understandable assumption, it might be problematic in practice for multiple reasons:

- (a) Collecting botnet metrics requires time and effort. How will the ISPs justify the business expense of this work, and the capital costs associated with instrumenting the parts of the ISP's network that may not currently be set up to gather the required data?
- (b) There are many ISPs in the United States. There are even more ISPs in other countries. Many ISPs will not participate. Incomplete participation (even simple addition and subtraction of participating ISPs) will complicate data interpretation and analysis.
- (c) ISPs have proven to be reluctant to share data about customer bot detections because of the distinct possibility that bot detection statistics will be misinterpreted. For example, if ISP A has a higher level of bot detections than ISP B, does that mean that ISP A is "better" at rooting out botnet customers than ISP B? Or does it mean that ISP A customers are inherently "less secure" than ISP B customers? Or does it mean that the bad guys are simply attacking ISP A customers more aggressively than ISP B?
- (d) Customers, third-party privacy organizations, and some governments (rightly or wrongly) may view ISP data sharing about botnets as a potential infringement of ISP customer privacy, even if all customer data is highly aggregated. (Notice the tension between customer privacy and the methodological ideal of gathering fine grained data with unique identifiers.)
- (e) Different ISPs measure botnet customers differently (efforts at standardization notwithstanding), undermining the comparability of inter-ISP botnet metrics. Different network-based detection techniques, and even identical techniques using different threat intelligence sources, have been shown to yield vastly different results.
- (f) Self-reported and unaudited bot data may be subject to error or manipulation, or at least the perception by cynics that it might not be fully candid/fully accurate/fully trustworthy.
- (g) Finally, we need to recognize that most bots are not domestic, while the Code participants are. Thus, US ISPs are poorly positioned to provide detailed botnet intelligence on most of the bots that are actually hitting US targets. We need other entities, entities with a global footprint, if you want consistent data with a global scope.

Entities With A Global Footprint: There are entities *other than ISPs* with consistent global visibility into the bot status of Internet systems and users:

- (a) Anti-virus or anti-spam companies with large international customer bases.

- (b) Operating system vendors that do periodic patching and malware removal (classic example: Microsoft with its Malicious Software Removal Tool that runs every month at patch time). Some of these entities already produce public reports about the data that they have collected. However, there are at least two important limitations to this data. Bots and other malware often attempt to disable operating system security measures and interfere with anti-virus protection, leading to a potential underreporting of infection levels. Secondly, these entities cannot report on every aspect of the Code, specifically notification.

Are we taking adequate advantage of data that has already been published? If not, why not? We believe that the FCC should not be "reinventing the wheel," particularly if there is no funding that can be used to ensure that botnet-related data is collected carefully and consistently. A list of some existing cyber security reports can be found in Appendix A to this guide.

VII. Sinkholing, DNS-Based Methods, Direct Data Collection and Simulations?

Sinkholing Specific Botnets: Sometimes a researcher or the authorities are able to gain control of part of a botnet's infrastructure. When that happens, the researcher or government person may be able to direct botnet traffic to a sinkhole, and use the resulting data visibility to measure certain aspects of a particular botnet.

Some might hope that sinkholing would provide a general purpose botnet estimation technique. Unfortunately, because this is a bot-by-bot approach, and requires the researcher or authorities to "inject" themselves into the botnet's infrastructure, it will not work to provide broad ISP-wide botnet measurements for all types of botnets. Many modern botnets now also take special care to prevent or deter efforts at sinkholing.

DNS-Based Approaches: Another approach that is sometimes mentioned is measuring botnets by their DNS traffic. That is, if you know that all botnet hosts "check in" to a specific fully qualified domain name, if an ISP sees a customer attempt to resolve that fully qualified domain name, there is a substantial likelihood that the customer is botnetted.³⁶

Big botnets might tend to make more extensive use of DNS than smaller and less sophisticated botnets, but caching and other subtleties associated with DNS can complicate DNS-based measurements, and of course, not all bots even use DNS.³⁷

Some might also be tempted to try an DNS Response Policy Zone (RPZ)-like approach (as implemented in BIND³⁸) to "take back" DNS and prevent bots from using DNS as part of their infrastructure. While this approach certainly has technical promise, any effort to instantiate policy via DNS is a potentially tricky one.³⁹

³⁶ One noteworthy exception: customers who are security researchers!

³⁷ Hypothetically, a botnet might choose to use raw IP addresses, or to use a peer-to-peer alternative to traditional DNS, such as distributed hash tables.

³⁸ <https://www.isc.org/software/bind>

³⁹ Remember the Internet's negative reaction to Stop Online Piracy Act (SOPA)/Protect IP Act (PIPA).

Directly Checking Systems to Find Botted Hosts? Assume that one wants a direct information-gathering approach that does not rely on ISPs providing data, or on third-party data sources. That is, one wants to go out and collect data directly, much as survey research groups survey entities about political viewpoints or consumer spending. How many individuals might one need to survey to get sufficient data about bottled users?

The required number of users will depend on the breakdown between bottled and non-bottled users, and the number of ISPs whose customers one would like to be able to individually track. If one has no "hints" about who may be a bottled user *a priori*, one must instead discover them at random. That may be daunting, at least if bots are indeed a "rare disease."

Let us arbitrarily assume we want 350 bottled users to study.

If 1.5% of all users are bottled, on average we would see 15 bottled users per thousand. Given that we want 350 bottled users, that would imply we would likely need to check $(350/15)*1,000=23,300$ users in order to find the 350 bottled users we needed.

But now, what if we now suppose that just 0.0711% of all users are bottled (recall that this was the CBL-reported rate for the US on December 9th, 2012)? On average we would see just 0.711 bottled users per thousand. To get 350 bottled users to study, we would likely need to check $(350/.711)*1,000=492,264$ users. That is a lot of data to collect.

Assume that we were charged with going out and checking 492,264 computers to see if those systems had been bottled. To keep this simple, let us assume that we will call a device bottled if a bot is found when we run a commercial antivirus product on that device.

If we assume that it would take an hour to run a commercial antivirus program on one machine (a very low estimate given the increasing size of consumer hard drives today), and technicians work 40 hours a week, it would take $492,264/40 = \sim 12,307$ "technician weeks" to scan 492,264 systems.

If a technician works 50 weeks a year, that would be 246.14 "technician years" worth of work.

If we assume an entry-level antivirus technician earns even \$50,000/year, and neglecting all other expenses (managerial/supervisory salary costs and software licensing costs and travel costs, etc.), our cost would be $246.14*50,000=\$12.3$ million.

Resistance to Government Scanning of Personal Computers: We suspect that many users would not be willing to allow a random government-dispatched technician to "scan their computer."

Personally owned computers often have intensely private files -- financial information (tax records, brokerage information, etc.); medical records; private email messages; goofy photographs, etc. In other cases, users may even have out-and-out illegal content on their systems (simple/common example: pirated software/movies/music).

Given these realities, many users would simply refuse to allow their device to be checked, even if they thought their device might actually be infected.

Volunteers? Some users who think that their systems might be infected might welcome the opportunity to have their systems scanned. Unfortunately, a "convenience sample" of that sort would not result in data that would allow us to generalize or extrapolate from the sample to the population as a whole.

Simulating Bot Infections in a Lab/Cyber Range? Another option, if we wanted to avoid the problems inherent in surveying/checking users, might be to try simulating bot infections in a lab or on a so-called "cyber range." While conceptually intriguing, this might not be easy. For example, the fidelity of the results from such a simulation would depend directly on researchers' ability to:

- (a) Replicate the full range of devices seen on the Internet (operating systems used, antivirus systems used, applications used, patching practices, etc. – do we have the data we would need to do that?).
- (b) Replicate the range of botnet malware seen on the Internet (constantly changing).
- (c) Accurately model the ISP response to the malware threat.

Given these difficulties, we believe this is a fundamentally impractical approach.

VIII. Recommendations

The CSRIC Working Group 7 recognizes that summary metrics with which to determine the effectiveness of the U.S. Anti-Bot Code of Conduct are not yet available.

The working group recommends that the following course be undertaken in order to establish such metrics for the future:

1. WG7 recommends that the FCC working in partnership with other federal government agencies and Industry facilitate the creation of case studies on bot mitigation activities, to examine metrics created around particular bot remediation efforts, by identifying up to three prolific bot-infection events occurring over the next year, and seek participation in the case studies from the multi-stakeholder community. Case Studies focus on specific events and provide a detailed analysis which can be used to better understand what is important to measure. These efforts will need to involve not only the ISP community but the larger Internet ecosystem as well.

A study currently in progress at a major US university in the south east to be published in March 2013 on ISP actions related to DNSChanger should be considered by the FCC as an initial case study. This effort, still in progress at the time of this writing, centers on the DNSChanger malware and related customer notification methods used by ISPs. This approach may well show the cooperative, collaborative steps required for the future development of overarching metrics involving multiple ISP approaches, and could be used as a model for future case studies. These steps are a microcosm of the recommendations in the Code.

2. We further recommend the FCC working in partnership with other federal government agencies and industry leverage industry sponsored pilot programs, e.g., M³AAWG Bot Metrics Pilot Program, to examine the collection and sharing of metrics around particular bot efforts, based on and incorporating results of the case studies, by identifying metrics to be collected and shared across the Internet ecosystem to help reduce bot infections. Pilot Programs differ from Case Studies in that they collect metrics on an on-going basis and provide participants with anonymized data which can baseline relative performance. Such programs are required to test which metrics may be useful and obtainable and which are not. Some metrics may, in fact, not be feasible to obtain at all. For comparative purposes, the metrics definitions must be reasonably standardized between ISPs. We anticipate that some metrics methods used by ISPs will lend themselves to comparative analysis and some will not. Participation in pilot programs will indicate which, if any, are viable.

3. We recommend that the FCC working in partnership with other federal government agencies and industry facilitate research in bot metric development by identifying gaps and shortfalls between outcome-based metrics needed to measure the effectiveness achieved by ISPs participating in the Code and metrics ISPs typically collect and share across the Internet ecosystem related to bot mitigation strategies. This research would seek to develop common definitions of bots for industry and methods to distinguish bots and bot families from other malware for measurement purposes. The research would provide input to industry efforts to define a set of implementable, outcome-based effectiveness metrics and standard methods of measurement based on best practices, case studies, and pilot programs.
4. We further recommend that the FCC working in partnership with other federal government agencies establish a vehicle such as a workshop or other education technique (i.e., a webcast) to foster ongoing dialogue around these issues. We encourage the inclusion of appropriate international participants (e.g., representatives from bot remediation bodies from Australia, Japan, etc.), as practical, in order to incorporate their learnings. WG7 believes this approach, incorporating and discussing the results of the above recommendations, will lead to further recommendations on Internet ecosystem multi-stakeholders approaches to best contain the spread of bot infections.

Metrics Guide Appendix A.

Examples of Data Driven Cybersecurity Reports

1. Composite Block List Statistics
<http://cbl.abuseat.org/statistics.html>
2. Kaspersky Security Bulletin/IT Threat Evolution
http://www.securelist.com/en/analysis/204792250/IT_Threat_Evolution_Q3_2012
3. McAfee's Quarterly Threats Report
http://www.mcafee.com/apps/view-all/publications.aspx?tf=mcafee_labs&sz=10®ion=us
4. Microsoft's Security Intelligence Report (SIR)
<http://www.microsoft.com/security/sir/>
5. Shadowserver Bot Counts
<http://www.shadowserver.org/wiki/pmwiki.php/Stats/BotCounts>
6. Symantec's Internet Security Threat Report (ISTR)
http://www.symantec.com/about/news/resources/press_kits/detail.jsp?pkid=threat_report_17

Metrics Guide Appendix B.

Another Data Collection Alternative, If Botnets Are A National Security Threat And Not Merely a Nuisance

While botnets are often thought of purely as a nuisance, e.g., as a source of spam and similar low grade unwanted Internet traffic, bots have also been used to attack government agencies and Internet-connected critical infrastructure. Viewed in that light, bots might properly be considered a threat to national security.

If bots are indeed a threat to national security, "other government agencies" may be able to directly apply "national technical means" to collect intelligence about botnets, including per-ISP estimates.

Such information, once collected, might then be able to be shared with appropriately cleared government officials with a legitimate need-to-know.

If domestic collection mechanisms are not an option or appropriate, it may also be possible to make estimates about domestic bot populations based on data collected by international partner agencies.

Appendix 5 – Metrics Glossary

Botnet Metrics Definitions

I. The Infection Lifecycle

1. Clean: For the purposes of the ISP voluntary Anti-Botnet Code of Conduct, a computer or other networked device will be considered "clean" when it (a) exhibits no externally discernable symptoms of infection (such as sending spam, participating in a distributed denial of service attack, or contacting a known command and control host), and/or (b) a review of the computer or other networked device with a generally accepted commercial or free/open source anti-virus program (using the most recently available definitions) finds no infection, and (c) the computer or other device otherwise appears to be operating normally in all respects. A newly purchased system shall be presumed to start in a clean status "out of the box," absent evidence to the contrary.

2. Vulnerable: A computer or other networked device shall be deemed "vulnerable" when it has one or more flaws or misconfigurations that render it potentially susceptible to compromise or infection. A common example of a vulnerable device is one that hasn't been patched, or which uses an easily guessed password for access. Note that a system may be "clean" yet "vulnerable" simultaneously, as in the case where a vulnerable system is protected by a compensating control (such as a firewall), thereby allowing the system to avoid becoming infected or compromised despite the presence of one or more vulnerabilities.

3. Infected: An infected computer or infected network device is one that has had malicious software or malicious firmware installed on/in it without knowing authorization. That malicious software or malicious firmware may be called a virus, a Trojan horse, a worm, a rootkit, a keystroke logger, a dialer, crimeware, spyware, adware, etc. A full definition of "malware" can be found in the glossary appearing in Appendix A to the FCC CSRIC "ABCs for ISPs."

4. Isolated: A system that is infected or compromised may be isolated to prevent it from generating unwanted Internet traffic. Isolated hosts are often put into "walled gardens" where they are limited to accessing a set of resources needed for remediation, or are allowed just to access life-safety services (such as VoIP telephony service for emergency use). There are different types of walled gardens with different restrictions to the Internet.

5. Offline: An offline system is one where no network access is allowed to/from that host whatsoever. Conceptually, think of an Ethernet connected host where the Ethernet cable has been disconnected (or the Ethernet switch port has been disabled), although obviously different technical processes are used in the case of cable modem connections, DSL connections, wireless access, modem access, etc.

6. Disinfected: A system shall be considered "disinfected" when, having been infected, it has been returned to a "clean" state (as previously defined above). The first step in disinfecting

a system is often running software to disable the malware and possibly clean the infected machine. More advanced malware may prevent the end user from updating their antivirus software and a tool must first be run to disable the malware before their security software can be reactivated or installed. In some cases, it may be necessary to format and reinstall the system from scratch to overcome particularly well-hidden persistent malware. Very advanced malware could infect computer BIOS or other embedded systems that might require replacing the computer.

7. Hardened: A hardened system is one that has been systematically configured so as to minimize exposure of the system's vulnerabilities (or potential vulnerabilities). This is known as reducing the attack surface. For example, among other things, a hardened system will be patched up to date, will have all unnecessary services disabled, will require encryption of all sensitive network traffic, will use strong passwords or multifactor authentication, will do secure logging to an off-system logging host, etc.

8. Reinfected: A system that has been disinfected but NOT hardened will often promptly become reinfected.

9. Compromised: While many vulnerable systems may be compromised as a result of being infected with malware, other vulnerable systems may be compromised as a result of weak passwords or misconfigurations (such as critical files that are unintentionally able to be modified by unauthorized parties). A compromised system is not trustworthy.

10. Managed: A so-called "managed host" is one that is centrally administered, rather than being self-administered by the system's user(s). Managed hosts are commonly seen in large corporations, and in government agencies.

11. Monitored: A monitored host is one that is continually (or at least periodically) scrutinized for things like anomalous network traffic or unauthorized changes to critical system files. Monitoring may take place via network security systems, such as Snort, or via host-based systems such as Tripwire.

12. Replaced: While most users will attempt to disinfect and harden an infected system, some may elect to replace that system with a new one instead. The prior system may then be sold or given to a third party, who may get the system along with any malware installed on it.

13. Shared: A shared system is one that's used by multiple individuals. A common example of a shared device would be a family device used by a parent or parents as well as by children or other family members. A shared device often seems to be more prone to infection (or other security issues) than a system that's used by only a single entity.

14. Orphaned: An orphaned device or orphaned program is an older one for which the vendor no longer releases even critical security/stability patches. Orphaned systems or programs generally cannot be hardened.

II. Ecosystem Roles

1. Customer: In the ABCs for ISPs context, the person who is paying an ISP for Internet service.
2. System Owner: The person who owns a given computer or other device.
3. System User: A person whom the system owner intentionally allows to use their computer or other device.
4. Support Person: For residential computers, a support person may be a family member or friend who helps the system owner or user to use and maintain their computer. A support person may also be a commercial computer support specialist hired for that purpose by the computer owner or user.
5. ISP Security/Abuse Team: The person or group at an Internet Service Provider who deals with complaints about a customer.
6. Vendor: The company that manufactured and marketed a computer system or software program. One might talk about an operating system vendor, an application software vendor, a hardware vendor, or an antivirus software vendor, for example.
7. Law Enforcement: A police officer, sheriff, federal agent, or other sworn individual with the power to investigate crimes, gather evidence and make arrests.
8. Regulator: A state or federal official tasked with managing business practices or other activity so as to ensure fundamental fairness or regulatory compliance. An example of a regulator would be the U.S. Federal Trade Commission. Regulators normally employ civil sanctions (such as administrative fines or civil lawsuits) rather than criminal sanctions (such as arrest/incarceration).
9. Unauthorized User: A person who uses a computer or other device without the intentional permission of the system owner, or someone who uses authorized access in excess of their authorization.
10. Malware Author: The programmer or programming team that designs and codes a piece of malware, such as a bot.
11. Botmaster: A botmaster is a person who operates a network of botted computers, often using them to send spam or attack other computers. The botmaster normally sends commands to "his" or "her" bots via a command and control host, e.g., a server under his or her control.
12. Affiliate: In this context, an affiliate is a person who helps to market a particular product or service in exchange for compensation, typically using pay-per-impression, pay-per-click, pay-per-install, or revenue sharing models:
 - (a) Pay-per-impression (PPI): affiliates are typically website owners who are paid according to the number of times a web site banner or other advertisement is shown to visitors

(b) Pay-per-click: in this model, affiliates are paid when someone actually clicks on an advertisement

(c) Pay-per-install: in this model, affiliates are paid when a program supplied to them is installed on a new system, either surreptitiously or with the knowing consent of the user (perhaps as part of a "sponsored access" offer for a program or site that would otherwise need to be purchased)

(d) Revenue sharing: in this model, affiliates are paid a percentage of the sales associated with the customers they refer.

13. List Seller: A list seller is someone who compiles and distributes lists of email addresses. For example, a spammer who wants to spamvertise an illegal online casino might purchase a list of email addresses known to be associated with online gamblers.

14. Bullet Proof Hosting Company: A so-called bullet proof hosting company is one that agrees to host a web site or other online presence notwithstanding complaints that may result from that activity, typically in exchange for the hosted party paying a premium price. Bullet proof hosting companies may be used to host spamvertised web sites, malicious software, child abuse materials, or other content likely to be unacceptable to regular hosting companies.

15. Bullet Proof Domain Name Registrars: A so-called bullet proof domain registrar is one that allows a spammer or other cyber criminal to register a domain name, and to keep that domain up, notwithstanding complaints that may be associated with that domain name. This service normally is provided for a premium over market domain name registration rates.

16. Payment Processor: When an affiliate makes a sale, payment is normally made by credit card. The entity that processes that credit card transaction is known as a "payment processor."

17. Drop Shipper: A drop shipper is an entity that manages order fulfillment for an affiliate program. For example, a drop shipper specializing in illegal pharmaceuticals may package and ship orders obtained by a pill spammer.

18. Online Currency Exchanger: Some affiliates may be paid using an online currency rather than via a mailed check or direct deposit. Online currency exchangers make it possible for some to purchase an online currency in exchange for cash, or vice versa.

19. Abuse Reporter: Third party reporting an abuse incident to an ISP, or through a clearinghouse (such as a computer security incident response team).

Appendix 6 - Related Industry Security and Metrics Activity

M³AAWG Bot Pilot Phase 1 Metrics

The mission of the M3AAWG Bot Metrics Pilot Program is to create a quorum of global ISPs who will anonymously submit data to permit factual sizing and locations of botnets and other malware around the world. The published metrics will serve as an objective tool for tracking the industry and governments' efforts at controlling the spread of botnets. Objectives of the program include: 1) To educate public policy makers about M3AAWG bot metrics to help frame the policy debate 2) To assess the efficacy of efforts to reduce bot infections. And 3) Provide a reliable view of the problem from those who own the data and see the issue first hand.

Basic rules of the program:

- Participation is voluntary
- Participants must be responsible for providing internet access
- Should commit to report metrics quarterly for 2 years, though a company may drop out if there are reporting problems
- Companies may be added anytime and should provide at least 2 quarters of reports for consistency
- All reports confidential; Executive Director receives reports and aggregates data
- Publication on M3AAWG website
 - After formal program approved by M3AAWG board of directors.

Total # of subscribers	This should be the number of subscribers for the business division being reported here. For example, you might not collect metrics from fiber links, etc.	# subscribers on last day of month
Infected Subscribers	A subscriber deemed infected one or more times within the reporting period counts as one infected subscriber. At this time we do not differentiate between reinfections and multiple different infections for the same subscriber. Definition of infected is left up to the provider, but generally needs to meet the bar for willing to notify the subscriber.	Count of <u>unique</u> subscribers with infections discovered in reporting period
Total number of subscribers notified	How many <u>unique</u> subscribers were notified of a problem by any means (SMS, phone call, email, web redirection/browser notification etc.) Multiple notices to the same subscriber count as one. This does not imply that the	# of unique infected subscribers in a reporting period who were notified

subscriber has read/received the notice.

Japan's Cyber Clean Center Metrics

Botnets are detected through the use of honeypots, which the Cyber Clean Center then analyze to identify the source of the infection. The CCC communicates the IP address and timestamp of a compromised machine to the participating ISP. And based on this information, the ISP may identify the customer and contact them by email, directing them to a CCC-sponsored webpage dedicated to botnet disinfection. Users may visit the site and download a tool to clean their device.⁴⁰ The Cyber Clean Center's objectives are to:

- reduce the number of bot-infected users,
- make bot removal tools, such as CCC Cleaner, widely available in Japan, and
- provide samples of malware to security vendors participating in the project.

In order to measure the effectiveness of the CCC in meeting their objectives, the CCC uses the following metrics:

- **Bot infection rate** defined as the ratio of infected broadband users to total broadband users.

Australia's iCode Metrics

Australia's ISP Voluntary Code of Practice, their iCode, is designed to provide a consistent approach for Australian ISPs to help inform, educate, and protect their customers in relation to cyber security risks.

- **Percentage of Compromised Internet User IP Addresses**⁴¹ defined as the ratio of infected subscribers divided by total subscribers.
- **iCode Adoption Rate**⁴² defined as the ratio of the number of subscribers of ISPs who have adopted the iCode to the total number of subscribers.
- **iCode ISPs**⁴³ defined as the total number of ISPs adopting the iCode.

Germany Anti-Botnet Initiative Metrics

Bot-infected machines are identified through the use of spam traps and honeypots operated by ISPs. Once an infected device is identified, ISPs inform their customers who are

⁴⁰ See https://www.ccc.go.jp/en_ccc/

⁴¹ See http://www.cso.com.au/article/427609/iaa_reviews_icode_cyber_security_war/#closeme

⁴² See http://www.nominum.com/assets/Documents/Webinar/Anti-botnet_lessons_learned_2012_05_10.pdf

⁴³ *Ibid.*

referred to a central help service (the Anti-Botnet Advisory Centre), which offers software tools for removing malware. Customers who are unsuccessful using these tools to recover their device may then contact their ISP for further assistance.

- **Anti-Botnet Initiative Adoption Rate**⁴⁴ defined as the ratio of the number of subscribers of ISPs who have adopted the German Anti-Botnet Initiative to the total number of subscribers.
- **Anti-Botnet Initiative ISPs**⁴⁵ defined as the total number of ISPs adopting the Anti-Botnet Initiative.

Ireland's Anti-Botnet Initiative Metrics

The Irish Reporting and Information Security Service (IRISS-CERT) identifies a device infected with malware or participating in a botnet and contacts the ISP. The ISPs contact end users through a variety of channels to inform them of the compromised machine. Customers are directed to the Irish Anti-Botnet Initiative website to download cleansing tools.

Finland's Anti-Botnet Initiative Metrics

The Finnish national CERT delivers lists of infected IP addresses to the affected ISP. The IP addresses are loaded into the ISPs' internally developed monitoring and alerting application. While applications are capable of automatically shutting down infected devices, the ISPs include additional steps to verify the external data before taking action to isolate infected devices.

Additional Bot Metrics Activities

Shadowserver Foundation⁴⁶ Metrics

A common problem with metrics is the lack of information about the exact methodology and corresponding time period of measurement. A positive example of how to handle this circumstance is the tracking information published by the non-profit organization, Shadowserver Foundation.⁴⁷

Shadowserver tracks various activities related to botnets. For example, regarding command-and-control servers, a constant number of between 5000 and 6000 online servers were observed in the course of the year 2010. For the purpose of counting infected devices, three differently parameterized metrics are published, representing a simplified definition of a botnet threat level. This metric is defined as a counter that decreases over time, so long as no

⁴⁴ See <http://www.oecd.org/internet/interneteconomy/45509383.pdf>

⁴⁵ *Ibid.*

⁴⁶ See <http://www.shadowserver.org/wiki/pmwiki.php/Information/Botnets>

⁴⁷ See <http://www.enisa.europa.eu/activities/Resilience-and-CIIP/critical-applications/botnets/botnets-10-tough-questions>

activity in the botnet is observed on a certain IP address. Every time the IP is active in the botnet, the counter is refreshed and set to its initial maximum value. Statistics are published for values of 5, 10, and 30 days respectively. A longer period clearly results in a higher bot count, as the chance increases that, during this timespan, malicious activity connected with an IP address can be observed.⁴⁸

Spamhaus CBL⁴⁹ Metrics

The Spamhaus Project is an international nonprofit organization whose mission is to track the Internet's spam operations and sources, to provide dependable realtime anti-spam protection for Internet networks, to work with Law Enforcement Agencies to identify and pursue spam gangs worldwide, and to lobby governments for effective anti-spam legislation.

⁴⁸ See <http://www.enisa.europa.eu/activities/Resilience-and-CIIP/critical-applications/botnets/botnets-10-tough-questions>

⁴⁹ See <http://www.spamhaus.org/organization/>