BITAG
Broadband Internet Technical Advisory Group
1550 Larimer Street, Suite 168
Denver, CO. 80202
Contact: Kaleb A. Sieh, Deputy Director, ksieh@bitag.org

BITAG Publishes Report on Port Blocking

Denver, CO (August 15, 2013): Today, the Broadband Internet Technical Advisory Group ("BITAG") announced the publication of its technical report on the subject of Port Blocking. The executive summary and recommendations of the report are attached here, and the full report itself can be found at: http://www.bitag.org/documents/Port-Blocking.pdf.

In the networking context, ports allow different applications running on an individual computer to share a single connection to the Internet. The practice of "port blocking" is often used, among other reasons, for security purposes to prevent unwanted or harmful traffic on a network and protect users. It is also capable of preventing the use of particular applications entirely, raising concerns that port blocking has the potential to be motivated by non-technical factors or otherwise construed as such.

BITAG's report aims to address some of these concerns by documenting how port blocking works, the rationales behind it, and its implications for different segments of the Internet ecosystem. The report also recommends best practices for entities that implement port blocking. Among other things, the report recommends that:

- ISPs should avoid port blocking unless they have no reasonable alternatives available for preventing unwanted traffic and protecting users. Further, if port blocking is deemed necessary, it should only be used for the purposes of protecting the implementing ISP's network and users.
- ISPs that can reasonably provide users opt-out provisions or exceptions to port blocking policies should do so.
- ISPs should publicly disclose their port blocking policies.
- ISPs should revisit their port blocking policies on a regular basis and reassess whether the threats that required the port blocking rules continue to be relevant.
- ISPs should make communications channels available for feedback about port blocking policies.
- Port blocking rules of consumer devices should be user-configurable.

Trace Hollifield, Director of Enterprise Security at Bright House Networks, and Jeffrey Swinton, Fellow at Verizon, were the lead editors of the report. Douglas Sicker, Executive Director of BITAG, Chair of BITAG's Technical Working Group and Endowed Professor of computer science at the University of Colorado Boulder, chaired the Port Blocking review. Kaleb Sieh, BITAG's Deputy Director, provided editorial assistance.

Questions or Comments? BITAG welcomes any questions, comments or suggestions. Please contact our Executive Director, Douglas Sicker, at dsicker@bitag.org or our Deputy Director, Kaleb Sieh, at ksieh@bitag.org.

– ATTACHMENT –

Executive Summary and Recommendations of BITAG Report on Port Blocking

The full report is available at: http://www.bitag.org/documents/Port-Blocking.pdf

Executive Summary

The term "port blocking" refers to the practice of an Internet Service Provider (ISP) identifying Internet traffic by the combination of port number and transport protocol, and blocking it entirely. Port blocking thus affects the traffic associated with a particular combination of port number and transport protocol on that ISP, regardless of source or destination IP address. The practice can potentially prevent the use of particular applications altogether by blocking the ports those applications use. Port blocks can be deployed in a range of network locations, from where the ISP connects with other networks to datacenters and customer locations.

The Internet was built around the premise of an open and shared environment. Additionally, Internet standards assume all hosts on the global Internet can connect directly to each other, on any specified port number. The practical reality is that blocking of Internet port numbers, either in the short or long term, is a technique that has been used by both wireline and wireless network providers for various reasons for over a decade.

One of the original and enduring motivations for blocking ports is to prevent network attacks and abuse associated with particular application protocols. Some network and security administrators view port blocking as a critical tool for securing systems and information, and see it as part of the ISP's mission to manage the security risk to its users from theft and destruction of personal information, business records, and other critical electronic forms of information. TCP port 25, used for sending email, is an example of a port that is blocked by some operators to prevent network abuse – such as spam email.

Port blocking has also been used to enforce ISPs' terms of service. Likewise, port blocking was once viewed as a useful tool for managing capacity and bandwidth-intensive applications such as peer-to-peer file-sharing applications on enterprise and university networks. However, increased network capacity and a variety of developments in the application space have caused most residential ISPs to seek other ways of managing capacity. Finally, though rare, port blocking has at times been used to hinder competing applications, such as Voice over IP (VoIP).

Port blocking is among a set of tools and tactics (Network Address Translation (NAT) being the other major example) that can compromise the original intent of ports: to provide reliable local addresses so that end systems can manage multiple communications at once.

Port blocking can complicate application design and development and create uncertainty about whether applications will function properly when they are deployed. Port blocking can also cause applications to not function properly or "break" by preventing applications from using the ports they were designed to use. One of the outcomes of port blocking is an increase in the use of "port overloading." Port overloading is a tactic whereby application developers will design applications to use a common port, in order to minimize the chance of a port blocking practice impacting the usability of that application.

Importantly, it may not be obvious to Internet users why an application affected by port blocking is not working properly, because the application may simply be unable to connect or fail silently. If error messages are provided, they may not contain specific details as to the cause of the problem. Users may seek assistance from the ISP's customer service, online documentation, or other knowledgeable sources if they cannot diagnose the problem themselves. The fact that the problem

could alternatively be caused by home networking equipment or a software-based port block complicates the process of diagnosis.

Users' ability to respond to port blocking depends on their technical sophistication and the extent to which workarounds are available. Overcoming port blocking may require the user to install a software update, change a configuration setting, request an opt-out from the ISP, or to upgrade their level of service (for example from residential to business). If these options are not available, or if users or customers lack the knowledge or willingness to pursue them, users may be prevented from using the blocked application altogether, or they may have to switch to a different application or a different network (from wireless to wireline, for example).

Because port blocking can affect how particular Internet applications function, its use has the potential to be anti-competitive, discriminatory, otherwise motivated by non-technical factors, or construed as such. As a result, the Broadband Internet Technical Advisory Group (BITAG) has a number of suggested practices when it comes to port blocking:

- **ISPs should avoid port blocking unless they have no reasonable alternatives available for preventing unwanted traffic and protecting users.** Further, if port blocking is deemed necessary, it should only be used for the purposes of protecting the implementing ISP's network and users. Port blocking should not be used for ongoing capacity management, to enforce non-security terms of service, or to disadvantage competing applications.

- **ISPs that can reasonably provide to their users opt-out provisions or exceptions to their port blocking policies should do so.** Whether opt-out provisions can be supported may depend on the particulars of the access network technology, the location port blocking is implemented in the network, administrative complexity, cost, and other factors.

- **ISPs should publicly disclose their port blocking policies.** The information should be readily available to both customers and non-customers alike, and should be as informative and concise as possible. For example, port blocking policies could be provided on the ISP's public facing website, on a page dedicated to summarizing or describing the respective ISP's network management practices.

  For persistent port blocks the information should include: (1) port numbers, (2) transport protocol (e.g., TCP or UDP), (3) the application(s) normally associated with the port(s), (4) the direction of the block – whether inbound or outbound, (5) a brief description of the reason(s) for the block, and (6) if opt-out provisions are available and how to request such.

- **ISPs should make communications channels available for feedback about port blocking policies.** Applications providers and consumers should have communications channels or other clear methods to discuss impacts caused by port blocking and to consider possible mitigations.

- **ISPs should revisit their port blocking policies on a regular basis and reassess whether the threats that required the port blocking rules continue to be relevant.** Some security threats are permanent and some are transitory or short-lived. Items such as spam prevention by blocking TCP port 25 from the customer are expected to last quite some time, while others such as blocks to prevent certain types of malicious software may be temporary.

- **Port blocking (or firewall) rules of consumers' devices should be user-configurable.** It is recommended that the documentation provided with each unit inform the consumer that port blocking or firewall rules have been implemented, which ports are blocked by default, and how consumers can modify those rules.