Competition and Consumer Protection in the 21st Century Hearings, Project Number P181201

Broadband Internet Technical Advisory Group (BITAG)
1550 Larimer Street, Suite 168
Denver, CO 80202

Contact email: info@bitag.org

BITAG is a multistakeholder group that brings together engineers and technical experts to develop consensus on broadband network management practices and related technical issues that can affect users' Internet experience (including impact to applications, content, and devices). This includes: educating policymakers on technical issues, addressing specific technical matters to help minimize related policy disputes, serving as a sounding board for new ideas and network management practices, identifying "best practices" of broadband providers and other entities, issuing advisory opinions, and otherwise providing technical guidance to industry and the public. BITAG has been in operation since the beginning of 2011. To date, BITAG has published nine reports providing a comprehensive understanding of broadband network management practices and Internet operations. BITAG is currently working on a report focused on Internet data collection and privacy.

The ultimate goal of BITAG is to bring Internet engineers and other technical experts together in "problem solving mode" to provide the best technical information and counsel to policymakers and the public – and to keep the focus on engineering principles first and foremost. BITAG facilitates conversations between relevant stakeholders affected by network management – and helps them work towards agreement concerning what is reasonable from a technical perspective. BITAG also works to address concerns on the part of policymakers or the public that a network management technique may be a pretext to discriminate or engage in anticompetitive conduct. Finally, BITAG publishes to the world the formal outcome of those conversations, in the form of technical reports, thus highlighting where there is agreement and where there may be disagreement among stakeholders.

BITAG reports are consensus documents with both educational and best-practice recommendation components, which are freely available to the public.

I. About BITAG – *An Overview*

Mission – BITAG brings together engineers and technical experts to develop consensus on broadband network management practices, with a focus on those practices where there is potential for differential impact or treatment, and thereby are open to claims of discrimination or anticompetitive conduct.

Technically Focused – Technical Working Group (TWG) participants must meet technical requirements.

Expeditious – The TWG operates under a 120-day "shot clock" within which it must analyze the technical topic and generate a report containing best practice recommendations.

Five Member Categories – BITAG has five participating member categories (Application Providers, Community Representatives, Content Producers, Equipment Manufacturers, Internet Connectivity Providers) and observers. Members provide funding for independent community representative engineers to represent those public interest groups that may not have technical staff of their own.

Balanced Processes and Consensus-based Decision-making – The TWG strives to operate on a consensus basis (as shown by the first nine reports), but also has weighted voting procedures in case of disagreement. All voting is weighted to reduce concerns with "forum packing," or in other words so that each member category has an equal say regardless of the number of members in the category. This weighted balance extends to the Board of Directors.

Work Product – BITAG has published nine (9) reports to date, which are outlined in detail in a separate section further below. Recent reports include: IoT privacy and security recommendations, differentiation (including prioritization), interconnection, VoIP impairment/failure/restriction, port blocking, real-time congestion management in Internet networks more generally. The most recent report on IoT security can be found at: http://www.bitag.org/report-internet-of-things-security-privacy-recommendations.php. BITAG's Technical Working Group is currently working on a report focused on Internet Data Collection and Privacy.

BITAG Leadership –

Executive Director & Chair of the Technical Working Group – Dr. Douglas C. Sicker is BITAG's Executive Director and Chair of the Technical Working Group. Doug is also currently the Department Head and professor of Engineering and Public Policy with a joint appointment in the School of Computer Science at Carnegie Mellon University, as well as the Director of CMU's CyLab Security and Privacy Institute. Previously, Doug was the DBC Endowed Professor in the Dept. of Computer Science at the University of Colorado at Boulder with a joint appointment in, and director of, the Interdisciplinary Telecommunications Program. Doug recently served as the Chief Technology Officer and Senior Advisor for Spectrum at the National Telecommunications and Information Administration (NTIA). Doug also served as the Chief Technology Officer of the Federal Communications Commission (FCC).

Deputy Director & General Counsel – Kaleb A. Sieh is BITAG's Deputy Director and General Counsel, where he performs an operations and legal role. Previously, he was a Research Fellow at the Silicon Flatirons Center at the University of Colorado Law School, where he engaged in research and

writing in telecommunications law, technology, and policy. He received his J.D. from CU-Law and his B.A. in Economics from CU as well.

Mission in more depth:

BITAG's mission is to bring together engineers and technical experts to develop consensus on broadband network management practices and related technical issues that can affect users' Internet experience (including impact to applications, content, and devices). This includes: educating policymakers on technical issues, addressing specific technical matters to help minimize related policy disputes, serving as a sounding board for new ideas and network management practices, identifying "best practices" of broadband providers and other entities, interpreting and applying "safe harbor" practices, issuing advisory opinions, and otherwise providing technical guidance to industry and the public.

The ultimate goal of BITAG is to bring Internet engineers and other technical experts together in "problem solving mode" to provide the best technical information and counsel to policymakers and the public – and to keep the focus on engineering principles first and foremost. To do so, BITAG focuses on those network management practices where there is potential for differential treatment or impact, and thereby are open to claims of discrimination or anti-competitive conduct. In this way, BITAG works to address concerns on the part of policymakers or the public that a network management technique may be a pretext to discriminate or engage in anticompetitive conduct.

BITAG then facilitates conversations between (ideally) all the relevant stakeholders affected by network management – and helps them work towards agreement concerning what is reasonable from a technical perspective. Finally, BITAG publishes to the world the formal outcome of those conversations, thus highlighting where there is agreement and where there may be disagreement among stakeholders.

Brief history:

BITAG was born out of a series of academic events and roundtable discussions starting in early 2008 at the Silicon Flatirons Center at the University of Colorado Law School in Boulder, Colorado.

BITAG began formation during the summer of 2010, with a broad cross section of stakeholders in the Internet ecosystem seeing a need for a group that could opine on reasonable network management techniques. Formal operations began in January 2011.

To date, BITAG has published nine (9) reports and held twenty-two (22) in-person multi-day meetings of its engineers and technical working group.

II. BITAG Reports – Exec Summaries and Recommendations

This section includes the executive summaries and recommendations from BITAG's already published reports. For full versions of each report please see www.bitag.org

A. Internet of Things (IoT) Security and Privacy Recommendations (Nov. 2016) – *Executive Summary, Observations, and Recommendations*

In the past few years, many of the new devices connected to the Internet have not been personal computers, but rather a variety of devices embedded with Internet connectivity and functions. This class of devices has generally been described as the *Internet of Things* (IoT) and has brought with it new security and privacy risks.

The term "IoT" has potentially broad scope. IoT can refer to deployments in homes, businesses, manufacturing facilities, transportation industries, and elsewhere. Thus, IoT can refer to much more than simply consumer-oriented devices. For the purposes of this report, we use the term IoT to refer solely to consumer-oriented devices and their associated local and remote software systems, though some or all of our recommendations may be more broadly applicable. This report is concerned with scenarios where consumers are installing, configuring, and administering devices that they lease or own.

The number and diversity of consumer IoT devices is growing rapidly; these devices offer many new applications for end users, and in the future will likely offer even more. Many IoT devices are either already available or are being developed for deployment in the near future, including:

- sensors to better understand patterns of daily life and monitor health
- monitors and controls for home functions, from locks to heating and water systems
- devices and appliances that anticipate a consumer's needs and can take action to address them (e.g., devices that monitor inventory and automatically re-order products for a consumer)

These devices typically interact with software running elsewhere on the network and often function autonomously, without requiring human intervention. In addition, when coupled with data analysis and machine learning, IoT devices may be able to take more proactive actions, reveal interesting and useful data patterns, or make suggestions to end users that may improve their health, environment, finances, and other aspects of their lives.

Although consumers face general security and privacy threats as a result of *any* Internet-connected device, the nature of consumer IoT is unique in that it can involve non-technical or uninterested consumers, challenging device discovery and inventory on consumer home networks as the number and variety of devices proliferate, impacts on the Internet access service of both the consumer and others that run on shared network links, and effects on other services in that when IoT

devices are compromised by malware they can become a platform for unwanted data traffic – such as spam and denial of service attacks – which can interfere with the provision of these other services.

Several recent reports have shown that some devices do not abide by rudimentary privacy and security best practices. In some cases, devices have been compromised and allowed unauthorized users to perform unauthorized surveillance and monitoring, gain unauthorized access or control, induce device or system failures, and disturb or harass authorized users or device owners.

Potential issues contributing to the lack of privacy and security best practices include: lack of IoT supply chain experience with security and privacy, lack of incentives to develop and deploy updates after the initial sale, difficulty of secure over-the-network software updates, devices with constrained or limited hardware resources (precluding certain basic or "common-sense" security measures), devices with constrained or limited user-interfaces (which if present, may have only minimal functionality), and devices with malware inserted during the manufacturing process.

The emergence of IoT presents opportunities for significant innovation, from smart homes to smart cities. In many cases, straightforward changes to device development, distribution, and maintenance processes can prevent the distribution of IoT devices that suffer from significant security and privacy issues. BITAG believes that following the guidelines outlined in this report may dramatically improve the security and privacy of IoT devices and minimize the costs associated with the collateral damage that would otherwise affect both end users and ISPs. In addition, unless the IoT device sector—the sector of the industry that manufactures and distributes these devices—improves device security and privacy, consumer backlash may impede the growth of the IoT marketplace and ultimately limit the promise IoT holds.

Observations. From the analysis made in this report and the combined experience of its members when it comes to Internet of Things devices, the BITAG Technical Working Group makes the following *observations*:

- **Security Vulnerabilities:** Some IoT devices ship "from the factory" with software that either is outdated or becomes outdated over time. Other IoT devices may ship with more current software, but vulnerabilities may be discovered in the future. Vulnerabilities that are discovered throughout a device's lifespan may make a device less secure over time unless it has a mechanism to subsequently update its software.

- **Insecure Communications:** Many of the security functions designed for more general-purpose computing devices are difficult to implement on IoT devices and a number of security flaws have been identified in the field, including unencrypted communications, data leaks from IoT devices.

- **Unauthenticated Communications:** Some IoT devices provide automatic software updates. Without authentication and encryption, however, this approach is insufficient, because the update mechanism could be compromised or disabled. In addition, many IoT devices do not use authentication in the course of communicating.
- **Unencrypted Communications:** Many IoT devices send some or all data in cleartext, rather than in an encrypted form. Communications in cleartext can be observed by other devices or by an attacker.
- **Lack of Mutual Authentication and Authorization:** A device that allows an unknown or unauthorized party to change its code or configuration, or to access its data, is a threat. The device can reveal that its owner is present or absent, facilitate the installation or operation of malware, or cause its core IoT function to be fundamentally compromised.
- **Lack of Network Isolation:** These devices also create new risks and are susceptible to attacks *inside* the home. Because many home networks do not, by default, isolate different parts of the network from each other, a network-connected device may be able to observe or exchange traffic with other devices on the same home network, thus making it possible for one device to observe or affect the behavior of unrelated devices.

- **Data Leaks:** IoT devices may leak private user data, both from the cloud (where data is stored) and between IoT devices themselves.
  - **Leaks from the Cloud:** Cloud services could experience a data breach due to an external attack or an insider threat. Additionally, if users rely on weak authentication or encryption methods for these cloud-hosted services, user data may also be compromised.
  - **Leaks from and Between Devices:** In some cases, devices on the same network or on neighboring networks may be able to observe data from other devices such as the names of people in a home, the precise geographic location of a home, or even the products that a consumer purchases.
- **Susceptibility to Malware Infection and Other Abuse:** Malware and other forms of abuse can disrupt IoT device operations, gain unauthorized access, or launch attacks.
- **Potential for Service Disruption:** The potential loss of availability or connectivity not only diminishes the functionality of IoT devices, but also may degrade the security of devices in some cases such as when an IoT device can no longer function without such connectivity (e.g., a home alarm system deactivating if connectivity is lost).

- **Potential That Device Security and Privacy Problems Will Persist:** IoT device security issues are likely to persist because many devices may never receive a software update, either because the manufacturer (or other party in the IoT supply chain, or IoT service provider) may not provide updates or because consumers may not apply the updates that are already available.
    - **Many IoT Devices Will Never Be Fixed:** Deploying software updates that patch critical security vulnerabilities is difficult in general. Many device vendors and manufacturers do not have systems or processes to deploy software updates to thousands of devices, and deploying over-the-network updates to devices that are operating in consumer homes is difficult, as updates can sometimes interrupt service and sometimes have the potential to "brick" the device, if done improperly. Additionally, some devices may not even be capable of software updates.
    - **Software Updates Address More Than Just Bugs:** Software updates are not simply intended to fix security or privacy bugs. They may also be intended to introduce major new functions, or improve performance and security.
    - **Consumers Are Unlikely to Update IoT Device Software:** Few end users consistently update device software of their own accord; it is best to assume that most end users will never take action on their own to update software.
- **Device Replacement May be an Alternative to Software Updates – for Inexpensive or "Disposable" Devices:** In some cases, replacing a device entirely may be an alternative to software updates. Certain IoT devices may be so inexpensive that updating software may be impractical or not cost-effective.

Recommendations. The BITAG Technical Working Group also has the following *recommendations*:

- **IoT Devices Should Use Best Current Software Practices:**
    - **IoT Devices Should Ship with Reasonably Current Software:** BITAG recommends that IoT devices should ship to customers or retail outlets with reasonably current software that does not contain severe, known vulnerabilities.
    - **IoT Devices Should Have a Mechanism for Automated, Secure Software Updates:** Software bugs should be minimized, but they are inevitable. Thus, it is critical for an IoT device to have a mechanism for automatic, secure software updates. BITAG recommends that manufacturers of IoT devices or IoT service providers should therefore design their devices and systems based on the assumption that new bugs and vulnerabilities will be discovered over time. They

should design systems and processes to ensure the automatic update of IoT device software, without requiring or expecting any type of user action or even user opt-in.

  - **IoT Devices Should Use Strong Authentication by Default:** BITAG recommends that IoT devices be secured by default (e.g. password protected) and not use common or easily guessable user names and passwords (e.g., "admin", "password").
  - **IoT Device Configurations Should Be Tested and Hardened:** Some IoT devices allow a user to customize the behavior of the device. BITAG recommends that manufacturers test the security of each device with a range of possible configurations, as opposed to simply the default configuration.

- **IoT Devices Should Follow Security & Cryptography Best Practices:** BITAG recommends that IoT device manufacturers secure communications using Transport Layer Security (TLS) or Lightweight Cryptography (LWC). If devices rely on a public key infrastructure (PKI), then an authorized entity must be able to revoke certificates when they become compromised, and manufacturers should take care to avoid encryption methods, protocols, and key sizes with known weaknesses. Additional encryption best practices include:
  - Encrypt Configuration (Command & Control) Communications By Default
  - Secure Communications To and From IoT Controllers
  - Encrypt Local Storage of Sensitive Data
  - Authenticate Communications, Software Changes, and Requests for Data
  - Use Unique Credentials for Each Device
  - Use Credentials That Can Be Updated
  - Close Unnecessary Ports and Disable Unnecessary Services
  - Use Libraries That Are Actively Maintained and Supported

- **IoT Devices Should Be Restrictive Rather Than Permissive in Communicating:** When possible, devices should not be reachable via inbound connections by default. IoT devices should not rely on the network firewall alone to restrict communication, as some communication between devices within the home may not traverse the firewall.

- **IoT Devices Should Continue to Function if Internet Connectivity is Disrupted:** BITAG recommends that an IoT device should be able to perform its primary function or functions (e.g., a light switch or a thermostat should

continue to function with manual controls), even if it is not connected to the Internet because Internet connectivity may be disrupted due to causes ranging from accidental misconfiguration or intentional attack. IoT devices that have implications for user safety should continue to function under disconnected operation to protect the safety of consumers.

- **IoT Devices Should Continue to Function If the Cloud Back-End Fails:** Many services that depend on or use a cloud back-end can continue to function, even if in a degraded or partially functional state, when connectivity to the cloud back-end is interrupted or the service itself fails.

- **IoT Devices Should Support Addressing and Naming Best Practices:** Many IoT devices may remain deployed for many years after they are installed. Supporting the latest protocols for addressing and naming will ensure that these devices remain functional for years to come.
    - **IPv6:** BITAG recommends that IoT devices support the most recent version of the Internet Protocol, IPv6.
    - **DNSSEC:** BITAG recommends that IoT devices support the use or validation of DNS Security Extensions (DNSSEC) when domain names are used.

- **IoT Devices Should Ship with a Privacy Policy That is Easy to Find & Understand:** BITAG recommends that IoT devices ship with a privacy policy, but that policy must be easy for a typical user to find and understand.

- **Disclose Rights to Remotely Decrease IoT Device Functionality:** BITAG recommends that if the functionality of an IoT device can be remotely decreased by a third party, such as by the manufacturer or IoT service provider, this possibility should be made clear to the user at the time of purchase.

- **The IoT Device Industry Should Consider an Industry Cybersecurity Program:** BITAG recommends that the IoT device industry or a related consumer electronics group consider the creation of an industry-backed program under which some kind of "Secure IoT Device" logo or notation could be carried on IoT retail packaging. An industry-backed set of best practices seems to be the most pragmatic means of balancing innovation in IoT against the security challenges associated with the fluid nature of cybersecurity, and avoiding the "checklist mentality" that can occur with certification processes.

- **The IoT Supply Chain Should Play Their Part In Addressing IoT Security and Privacy Issues:** End users of IoT devices depend upon the IoT supply

chain, from manufacturer to retailer, to protect their security and privacy, and some or all parts of that IoT supply chain play a critical role throughout the entire lifecycle of the product. In addition to other recommendations in this section, BITAG recommends that the IoT supply chain takes the following steps:

- **Privacy Policy:** Devices should have a privacy policy that is clear and understandable, particularly where a device is sold in conjunction with an ongoing service.
- **Reset Mechanism:** Devices should have a reset mechanism for IoT devices that clears all configuration for use when a consumer returns or resells the device. The device manufacturers should also provide a mechanism to delete or reset any data that the respective device stores in the cloud.
- **Bug Reporting System:** Manufacturers should provide a bug reporting system with a well-defined bug submission mechanisms and documented response policy.
- **Secure Software Supply Chain:** Manufacturers should protect the secure software supply chain to prevent introduction of malware during the manufacturing process; vendors and manufacturers should take appropriate measures to secure their software supply chain.
- **Support IoT Device for Entire Lifespan:** Manufacturers should support for an IoT device throughout the course of its lifespan, from design to the time when a device is retired, including transparency about the timespan over which they plan to provide continued support for a device, and what the consumer should expect from the device's function at the end of the device's lifespan.
- **Clear Contact Methods:** Manufacturers should provide clear methods for consumers to determine who they can contact for support and methods to contact consumers to disseminate information about software vulnerabilities or other issues.
- **Report Discovery and Remediation of Vulnerabilities:** Manufacturers should report discovery and remediation of software vulnerabilities that pose security or privacy threats to consumers.
- **Clear Vulnerability Reporting Process:** Manufacturers should provide a vulnerability reporting process with a well-defined, easy-to-locate, and secure vulnerability reporting form, as well as a documented response policy.

B. Differentiated Treatment of Internet Traffic (Oct. 2015) – *Executive Summary, Observations, and Recommendations*

The Internet is composed of interconnected networks, each having its own architecture and technical characteristics. The data transmitted across these networks is formatted as packets containing information payloads encapsulated within one or more headers, which in turn provide the information needed by networks to deliver the packets to their destinations. As these packets travel across networks, they contend with other packets for network resources. Contention can occur at any point where two or more packets can compete for a resource at the same time. The simplest way to handle such requests would be on a first come, first served basis (also known as First In First Out, or FIFO). In practice, however, network operators make many exceptions to FIFO, using the packet header information to classify packets into flows and treating those flows differently, for example rearranging the order or the timing with which packets are sent, or sending them along different network paths.

Differentiated treatment of Internet Access Service traffic has been a subject of debate and regulatory scrutiny. In February 2015, the Federal Communications Commission (FCC) adopted Open Internet rules that address paid prioritization as well as other topics [1]. This report touches on a broad range of questions associated with differentiation, but is not intended to address or analyze the economic, legal, regulatory, or public policy issues that the differentiated treatment of Internet access service traffic may raise, focusing instead on the technical issues. The ability to treat traffic differentially has been built into Internet protocols from the beginning. The specifications for both IPv4 and IPv6 have included fields to support traffic differentiation since their inception (initially IPv4's Type of Service or ToS field) to indicate to routers the quality of service desired, in terms of queuing precedence and routing parameters around delay, rate, and reliability. This was changed to more generic service descriptions with the definition of the Differentiated Services Field, and implemented in IPv4 and IPv6. Notably, traffic differentiation in this sense has not been implemented in multi-provider environments, although it is extensively used within specific networks. End to end deployment would require the harmonization and cooperation of a large number, if not all, of the relevant network operators.

In its broadest sense, traffic differentiation includes any technique that classifies and applies potentially different treatment to two or more traffic flows contending for resources on a network (a flow being a group of packets that share a common set of properties). Differentiated treatment of network traffic is a two-part process: (1) traffic is classified into traffic streams, and (2) a prescribed set of actions is applied to each stream. This treatment may determine the order in which routers and switches send packets from different flows across the link, the rate of transmission of a given flow, or even whether certain packets are sent at all.
While the techniques used for traffic differentiation overlap with those used to manage congestion, differentiation has a broader purpose that includes meeting

service level agreement (SLA) guarantees and selecting paths for traffic from different applications, among other things. Differentiated treatment of traffic can also contribute both to the efficiency of a network and to the predictability of the manner in which network resources are shared.

Differentiation can be complex, and a common vocabulary is key. This report uses the terms "differentiated treatment" or "differentiation," as opposed to "prioritization" when referring to the full range of treatments that may be applied to traffic flows. The technical definition of "prioritization" is narrow and generally applies only to certain scheduling, dropping, and marking techniques. This report uses "differentiation" in a much broader sense, including most of the ways in which packets may be treated differently from each other while en route to their respective destinations across one or more networks. The scope of differentiation in this report encompasses the classic techniques of scheduling, shaping and queue management by which packets are processed at a network node, and also includes the techniques by which traffic flows are segregated or forwarded onto different physical or logical network paths where they may encounter greater or lesser propagation delays or contention for resources.

This report addresses differentiation applied to traffic on Internet access services, as well as the impacts to Internet access services when differentiation is applied to other traffic carried over the same network. Traffic for mass-market Internet access services is often carried over a common infrastructure with traffic associated with other IP services, as well as the network management traffic used to control devices and report status from them. Since differential treatment of other network traffic has the potential to affect the performance of Internet access services, it is considered here.

The subjective experience perceived by the user of a networked application is known as Quality of Experience, or QoE, and the factors that contribute to QoE vary significantly from one application to the next. In contrast, Quality of Service, or QoS, describes the performance of a network service using objective metrics such as throughput, delay, delay variation, and loss. The relationship between QoS and QoE is highly dependent on the type of application, but variations in QoS have been mapped to corresponding variations in QoE for a number of applications. It is possible to use knowledge about the relationships between network performance parameters and their effects on QoE to attempt to optimize the performance of network flows for their intended applications. Differentiation is often also used to address impairments to QoS.

Broadband networks use different network architectures and access technologies. Several of these network architectures have developed to take advantage of existing access infrastructure that was originally deployed for other services – for example, telephone service over twisted copper pairs or video over coaxial cable. Other networks were developed to meet specific needs, such as for mobility or for access in remote rural areas. In many cases, differences in network design can be traced to

the different characteristics of the access technology used. Access technologies can require different approaches to differentiation of traffic.

Observations. From the analysis made in this report and the combined experience of its members when it comes to the differentiated treatment of Internet traffic, the BITAG Technical Working Group makes the following *observations*:

- **TCP causes recurring momentary congestion.**

  When TCP transfers a large file, such as video content or a large web page, it practically guarantees that it will create recurring momentary congestion at some point in its network path. This effect exists by design, and it cannot necessarily be eliminated by increasing capacity. Given the same traffic load, however, the severity of the momentary congestion should decrease with increased capacity.

- **A nominal level of packet discard is normal.**

  Packet discard occurs by design in the Internet. Protocols such as TCP use packet discard as a means of detecting congestion, responding by reducing the amount of data outstanding and with it self-induced congestion on the transmission path. Rather than being an impairment, packet discard serves as an important signaling mechanism that keeps congestion in check.

- **The absence of differentiation does not imply comparable behavior among applications.**

  In the absence of differentiation, the underlying protocols used on the Internet do not necessarily give each application comparable bandwidth. For example:

  - TCP tends to share available capacity (although not necessarily equally) between competing connections. However, some applications use many connections at once while other applications only use one connection.
  - Some applications using RTP/UDP or other transport protocols balance transmission rate against experienced loss and latency, reducing the capacity available to competing applications.

- **Differentiated treatment can produce a net improvement in Quality of Experience (QoE).**

  When differentiated treatment is applied with an awareness of the requirements for different types of traffic, it becomes possible to create a benefit without an offsetting loss. For example, some differentiation techniques improve the performance or quality of experience (QoE) for particular applications or classes of applications without negatively

impacting the QoE for other applications or classes of applications. The use and development of these techniques has value.

- **Access technologies differ in their capabilities and characteristics.**

  Specific architectures and access technologies have unique characteristics which are addressed using different techniques for differentiated treatment.

- **Security of traffic has at times been downgraded to facilitate differentiation techniques.**

  Encrypted traffic is on the rise and it has implications for current differentiation techniques. In response to this increase, some satellite and in-flight network operators have deployed differentiation mechanisms that downgrade security properties of some connections to accomplish differentiation. The resulting risks to the security and privacy of end users can be significant, and differentiation via observable information such as ports and traffic heuristics is more compatible with security.

Recommendations. The BITAG Technical Working Group also has the following *recommendations*:

- **Network operators should disclose information on differential treatment of traffic.**

  In previous reports, BITAG has recommended transparency with respect to a number of aspects of network management. BITAG continues to recommend transparency when it comes to the practices used to implement the differential treatment of Internet traffic.

  Specifically with respect to consumer-facing services such as mass-market Internet access, network operators should disclose the use of traffic differentiation practices that impact an end user's Internet access service. The disclosure should be readily accessible to the public (e.g. via a webpage) and describe the practice with its impact to end users and expected benefits in terms meaningful to end users. The disclosure should include any differentiation amongst Internet traffic and should disclose the extent and manner in which other services offered over the same end user access facilities (for example video services) may affect the performance of the Internet access service.

- **Network operators and ASPs should be encouraged to implement efficient and adaptive network resource management practices.**

  In a previous report BITAG recommended that ASPs and CDNs implement efficient and adaptive network resource management practices; we reiterate that recommendation here, extending it to network operators. Examples of such practices might target the minimization of latency and variation in latency induced in network equipment, ensuring sufficient bandwidth for expected traffic loads, and the use of queue management techniques to manage resource contention issues.

- **Quality of Service metrics should be interpreted in the context of Quality of Experience.**

  Common Quality of Service metrics, often included in commercial service level agreements, include capacity, delay, delay variation, and loss rate, among other things. From the viewpoint of the end user application, these metrics trade off against each other and must be considered in the context of Quality of Experience. For example, since TCP Congestion Control and adaptive codecs depend on loss to infer network behavior, actively trying to reduce loss to zero leads to unintended consequences. On the other hand, non-negligible loss rates often directly reduce the user's Quality of Experience. Hence, such metrics should be interpreted in the context of improving user experience.

- **Network operators should not downgrade, interfere with, or block user-selected security in order to apply differentiated treatment.**

  Network operators should refrain from preventing users from applying over-the-top encryption or other security mechanisms without user knowledge and consent. Networks should not interfere with, modify, or drop security parameters requested by an endpoint to apply differentiated treatment. Given the potential for possible exposure of sensitive, confidential, and proprietary information, prior notice should be given to end users of traffic differentiation features that affect security properties transmitted by endpoints.

C. Interconnection (Nov. 2014) – *Executive Summary*

The term "interconnection" refers to the various means by which network providers attach to and move traffic between one another, and is a collection of business practices and technical mechanisms that allow individually managed networks to connect together for this purpose. There is no central authority that manages Internet interconnection – the overall system arises because of the many bilateral and multilateral decisions that various actors make to interconnect. In contrast to

the telephone system, where interconnection is performed in a highly regulated environment, interconnection in the Internet (in most parts of the world) remains a private-sector matter.

The topic of Internet interconnection is receiving increased attention as the Internet ecosystem continues to evolve. Networks of all types interconnect among one another, including those of Internet access providers, content providers, academic institutions, and commercial enterprises. Internet connectivity is achieved by passing pieces of data, called packets, from a connected device through networking equipment, known as routers, operated by one or more network providers until those packets are delivered to the desired destination. The mechanisms that implement interconnection thus serve both technical and business purposes, and discussion concerning the technology of interconnection must, of necessity, refer to business issues to some extent – as many of the mechanisms can only be understood in that context.

Network interconnection in the United States has evolved significantly since the early days of the Internet, and today is a complex amalgam of models incorporating new connectivity options, delivery options, traffic management requirements and business practices. It is important to note the difference between the two dominant forms of interconnection, which are: (1) *transit* – where access to every publicly reachable destination on the Internet is provided for a fee; and (2) *peering* – where customer traffic is exchanged between two networks and the access provided is only to each other's network and customers. Further, when two networks peer there can also be both "settlement free" (without requiring payment) and paid arrangements.

Network operators are motivated to peer for a variety of reasons that may include both business and technical motivations. Each network operator stipulates the technical and operational criteria used to evaluate what networks they will interconnect with, and many of these requirements are made publicly available online. Connecting networks does not come without costs, and a decision to interconnect requires careful consideration of the benefits compared to the costs incurred to connect at each location.

The two common options for interconnection are either through a private bilateral arrangement between two networks using a dedicated physical connection (called a "direct interconnection"), or a multilateral arrangement where all networks connect into a public Internet Exchange switch. An Internet Exchange is a service that uses a switch infrastructure (sometimes referred to as a switch fabric or peering exchange) to provide connectivity between multiple networks. Interconnection methods are constantly evolving, and one of the more important developments in interconnection is the use of content delivery networks (CDNs). CDNs provide a more efficient means of distributing content by placing content and applications on servers distributed closer to, and sometimes within, the destination network – essentially bringing data (e.g., popular content) closer to the requestor instead of

delivering the data across a potentially significant portion of the Internet. The introduction of CDNs and IXs has contributed to the "flattening" of the historic hierarchical model of Internet interconnection.

Internet traffic has grown rapidly since the Internet's inception, and this has often been driven by the growth of popular applications. Managing the exchange of Internet traffic between networks is accomplished primarily through the use of an inter-network routing protocol called the Boarder Gateway Protocol (BGP). BGP offers network administrators the ability to implement routing policy, or in other words how traffic flows through a network. BGP's design offers limited support for inbound (traffic destined into one's network) traffic control.

Peering connections are increasingly the primary interconnection paths between networks, supplanting the model of hierarchical interconnection via a small group of long-distance network providers. In most cases, two parties seeking to interconnect are able to come to terms. In some cases after an agreement is reached, however, traffic volumes or other factors may change, which in rare cases may lead to "de-peering" events if the changes are significant enough. More commonly, such changes lead to a renegotiation of the manner or type of interconnection agreement between the two parties. Although peering disputes over traffic imbalances, and other reasons, are not new, peering disputes in the U.S. have been increasingly publicized in recent years.

In some cases traffic can flow contrary to the intentions of network operators, either in error or due to malicious activity. There are a number of important security considerations when connecting two networks. There are numerous types of attacks, as well as various motivations that may drive attackers. There are also a number of potential mitigations, as well as efforts to make routing more secure through new routing protocol extensions, notably RPKI and BGPSec.

This report provides a reference on the subject of Internet interconnection, and presents a detailed review on how networks connect, the development and changes in connection models, motivations for connection, how networks manage traffic between each other and some of the challenges that arise as networks evolve.

D. VoIP Impairment, Failure, and Restrictions (June 2014) – *Executive Summary and Recommendations*

IP networks have supported voice communications for some time. Voice over IP (VoIP) services allow users to make calls between IP-based endpoints and to interconnect with the traditional public switched telephone network. VoIP applications use a variety of methods and protocols to manage connections and exchange media (i.e., the content of voice or video communications) over IP-based

networks such as the Internet. In the majority of VoIP applications, connections are managed or controlled using one protocol or set of protocols, and the media is exchanged among the parties involved in the connection using a different transport protocol or set of protocols. The purpose of these separate connection control and signaling protocols is to allow the parties involved in the communication to establish, control, and terminate connections. Many services that provide VoIP capability also support other forms of real-time media (video or screen sharing, for example).

This report uses the term *VoIP impairment* to refer to anything that prevents a VoIP application from being used in the manner desired by a user. An impairment affecting a VoIP service can occur anywhere along the data path, including in the end devices. For example, impairment may be the result of actions by a network operator, the VoIP provider itself, or the provider of a smartphone's operating system. Impairment could likewise be the result of a poorly implemented network device or application, or may occur as a result of a configuration or misconfiguration of a home network. VoIP services can be rendered unusable if their quality is sufficiently reduced to prevent meaningful audio exchange between the participants, even if some VoIP traffic is still exchanged.

VoIP failure, as defined in this report, encompasses particular kinds of VoIP impairments that arise when VoIP calls cannot be established at all, or when no media is capable of being exchanged between VoIP endpoints. When network operators, VoIP providers, operating system vendors, or application store providers take steps that cause VoIP failures or prevent VoIP from being used, those steps are considered *VoIP restrictions* for the purposes of this report. The term *VoIP impairments* (plural) is used as shorthand for *VoIP impairment, failure, or restrictions* in this report. Issues related to interoperability between different VoIP services are out of scope.

This report discusses: (1) how VoIP works; (2) how VoIP may be impaired or restricted; (3) methods for mitigating VoIP impairments; and (4) recommendations concerning VoIP impairment, failure, and restrictions. The report focuses on VoIP impairments that may occur in residential or mobile networks. This report makes no assumptions concerning the motivations behind actions that result in VoIP impairment, failure, or restrictions, or about the frequency or scale of such occurrences.

There are a variety of technical causes of VoIP impairment, failure, and restrictions:

Port blocking. In the architecture of the Internet, communication between two systems is identified by five fields: (1) the source IP address, (2) the destination IP address, (3) the transport protocol in use, (4) the source port, and (5) the destination port used by the transport protocol. The pair of IP addresses representing two systems identifies all of the communication sessions between them, whereas the port number pair characterizes an individual communication

session between the two systems. If traffic is prevented from flowing to or from the particular ports used by VoIP applications, VoIP failure can occur. This practice is known as port blocking.

Application-Level Gateways (ALGs). Many endpoints on the Internet sit behind a Network Address Translation (NAT) device. NAT allows multiple end devices within a local network (the network in a single home, for example) to share a single public IP address. NAT is challenging for applications such as VoIP that require reachability from the Internet. A variety of mechanisms have been developed for overcoming these challenges, including the use of Application-Level Gateways (ALGs), which can automatically detect traffic associated with particular VoIP services and help that traffic pass through NAT devices. However, ALGs may also affect traffic handling in such a manner as to impair the correct operation of other VoIP services.

Other network-based causes. The use of some network-based techniques that attempt to identify VoIP requests or VoIP traffic may result in VoIP impairment. These techniques may leverage the Domain Name System (DNS) infrastructure, deep packet inspection, or network-based flow policing and filtering that attempt to identify patterns of likely VoIP traffic.

Restrictions based on device, application, or application store. Actions taken on a device, in a VoIP application, or in an application store may restrict the use or functionality of VoIP applications. These examples often reflect business arrangements or agreements between application store providers, operating system vendors, device manufacturers, and/or mobile carriers.

When VoIP impairments occur, it may be very difficult for a user to solve the problem, even for technically sophisticated users. Whether a mitigation or workaround solution exists, and how difficult it may be to implement, depends on the mechanism that is impairing the VoIP connection. If mitigation options are unavailable, or if users lack the knowledge or willingness to pursue such, those users may be prevented from using VoIP altogether, or may need to switch to a different VoIP application. Mitigations available to application providers for dealing with VoIP impairment also depend on how the impairments are effectuated.

VoIP impairment, failure, and restrictions can create difficulties for VoIP users and may deter adoption of over-the-top VoIP services. VoIP impairments can also create difficulties for the operators of VoIP services and providers of VoIP applications, who may need to troubleshoot or work around impairments (where possible) to enable or restore their users' connectivity.

BITAG's Technical Working Group recommends the following to minimize the occurrences and impact of VoIP impairment, failure, and restrictions:

- **Network operators should avoid impairing or restricting VoIP applications unless no reasonable alternatives are available to resolve**

technical issues. Certain network management actions may have the effect of limiting or restricting VoIP traffic as a method of ensuring network integrity. Examples include port blocks or traffic limitations implemented when a customer uses a vulnerable VoIP service that can be exploited by attackers for the purpose of flooding the network with unwanted traffic. In adopting any approach that has the effect of limiting the use of VoIP, network operators should seek to minimize the impact of the approach on legitimate VoIP use.

- **VoIP-related ALGs in operator-supplied home routers should minimize their impact on traffic other than the operator's VoIP service where possible.** VoIP-related ALGs can interfere with some VoIP services while attempting to facilitate NAT traversal for other VoIP services. Because of these problems, BITAG recommends that VoIP-related ALGs in operator-supplied home routers should either allow the VoIP-related ALGs to be disabled for customers who do not subscribe to the operator's VoIP service or minimize or avoid impact to independent VoIP services and all other traffic not associated with the operator's own VoIP service. Where possible, VoIP-related ALGs in operator-supplied home routers should be disabled by default. ALGs for real-time applications (including VoIP) can be problematic for services other than VoIP, but recommendations concerning ALGs more broadly are outside the scope of this report.

- **Manufacturers of home routers should disable VoIP-related ALGs by default.** Some consumers purchase their home routers from retailers rather than from network operators. To limit the impact of VoIP-related ALGs on VoIP services, home routers sold to consumers should have VoIP-related ALGs disabled by default.

- **Port blocking rules in consumer equipment should be user-configurable.** The port blocking (or firewall) rules of consumers' home routers should be user-configurable, whether the routers are provided by the ISP or purchased separately by the consumer. By making these rules user-configurable, technically sophisticated users may be able to eliminate port blocks that prevent them from using VoIP services. It is recommended that the documentation provided with the consumer equipment inform the consumer that port blocking or firewall rules have been implemented, the default ports blocked, and how consumers can modify those rules.

- **If network operators intentionally use network policies or practices that impair or restrict VoIP, they should provide disclosures about those policies and practices and provide communications channels for feedback.** BITAG recommends that network operators disclose their policies and practices that may or could result in VoIP impairment, failure, or restrictions. The information should be readily available to both customers

and non-customers alike. For example, such policies could be provided on the operator's public-facing web site or on a page dedicated to summarizing or describing the ISP's network management practices. If specific VoIP applications are impaired or restricted, those applications should be listed by name, along with a brief description of the reason for the impairment or restriction. BITAG also recommends that ISPs provide a communications channel or other clear method for application providers and consumers to discuss the impact of VoIP impairment, failure, and restrictions, and possible mitigations.

- **Application developers should design VoIP applications to be port-agile where possible.** BITAG recommends that VoIP application developers design VoIP applications and services to be port-agile where possible. Applications designed to tolerate random source ports or to allow port selection to be user-configurable are better able to avoid VoIP impairments that result from port blocking or contention between multiple services for the same port. Whether particular applications can be re-designed to be port agile may depend on whether re-designed versions of the application can be made compatible with existing versions or other existing applications.

E. Congestion Management Report (Oct. 2013) – *Executive Summary and Recommendations*

The Internet, as is the case with many other networks such as highways and electricity grids, operates under the assumption that capacity will be set to a level such that total peak demand will occasionally exceed capacity. Further, the Internet is designed so that multiple users may dynamically share capacity and multiple services may share the same network links and routers, which is more efficient than offering individual users dedicated capacity or different services using separate links and routers.

Every link and router in the various networks that make up the Internet has a limit on its capacity to handle data. The capacity of each link and router in individual networks is determined by the equipment installed by the entity that runs each network in an attempt to optimize performance and cost; the lower the capacity relative to expected demand, the greater the probability that demand upon that link or router at times may exceed its capacity.

Significantly, a user's instantaneous demand for broadband Internet is bursty, meaning that it changes rapidly in time – and when aggregate instantaneous demand exceeds capacity on a network it causes congestion, which can degrade performance.

Network operators typically estimate demand months to years in advance, and use such demand estimates to plan a schedule for capacity upgrades. Since it may take months to implement a capacity upgrade, the time scale for managing congestion in this manner is months to years. Thus, although capacity planning can greatly affect how much congestion occurs on a network over time, it cannot react to congestion as it occurs.

The impact of congestion upon applications depends on the duration of congestion – which can vary from thousandths of a second up to hours or more – and the nature and design of the application. If the duration of congestion is short enough or the application is tolerant enough of congestion, a user will not notice any degradation in performance. Congestion is thus a problem only when its duration is long enough to be disruptive to applications. Congestion in a network can occur for a wide variety of reasons, some of which can be anticipated and some of which cannot.

This report describes how network resources are allocated on a short time scale in order to, among other objectives, manage congestion on the network, and how such congestion management impacts applications and users.

Congestion management practices are an important subset of network management practices implemented by a variety of parties or organizations, including Internet Service Providers (ISPs) and Application Service Providers (ASPs). Policymakers have expressed great interest in learning what congestion management practices are used in the Internet and how these practices impact users and the broader Internet ecosystem. Furthermore, an understanding of congestion management techniques and practices is crucial in discussions about reasonable network management.

One of the key design questions about any congestion management practice concerns the subset of network traffic to which the practice is applied, and its impact upon users and applications. Network operators apply some practices to all traffic on their networks, whereas in other cases practices are applied only to the traffic of specific users or to the traffic associated with specific applications. Application- or user- based congestion management practices may achieve better performance for selected applications. They also may enable service providers to offer connectivity products that cater to particular customer's tastes or needs. However, they add complexity, which may result in added costs that each network operator will evaluate. In some cases application- or user- based congestion management practices may be harmful to applications.

Congestion management practices are composed of generic technical building blocks, described in this report as traffic management "techniques". This report discusses a range of user- and application- based congestion management techniques, including classification of packets, reservation of resources for particular network flows, storage of content in multiple locations, rate control, routing and traffic engineering, packet dropping, and packet scheduling.

Congestion management techniques may be combined to offer a collection of capabilities in various network architectures, and can create services with differentiated performance either within an operator's network or end-to-end. There are also architecture-specific implementations of congestion management techniques for broadband Internet access over cable, telephone, and cellular networks and for Content Delivery Networks. The offerings of a service provider often include multiple services that may utilize the same network links and routers. While there are benefits and efficiencies to sharing capacity between multiple services, such sharing of capacity also requires the use of congestion management practices.

Congestion management "practices" are the uses of particular techniques by particular network operators to avoid, limit, or manage congestion. This report illustrates a range of congestion management practices that show how providers may combine user- or application- based congestion management techniques, including traffic shaping, prioritization, transcoding, resource reservation, and preferential treatment.

The report begins in Sections 1 and 2 by giving an overview of congestion and BITAG's interest in the issue. Section 3 defines congestion and describes instances in which congestion can occur, the locations in the network where congestion can occur, the indicators of congestion, and the impact congestion can have on applications.

In Section 4, the report articulates the differences between congestion management techniques and congestion management practices, and describes the different time scales at which congestion can be seen to occur in the network. This section also describes the parties that implement congestion management practices and on what basis.

Although all congestion management is important, in order to limit scope and length Sections 5-7 focus on congestion management techniques and practices that: (1) are implemented or potentially implemented in a network that supports consumer broadband Internet access services; (2) act on a time scale of minutes or less; (3) are used for purposes of congestion management; and (4) are based on user or application.

In Section 5, the report focuses on specific congestion management techniques. Section 6 gives specific examples of congestion management practices that are based on user or application. Finally, Section 7 gives the Technical Working Group's recommendations.

At a high level, the recommendations of BITAG's Technical Working Group user are:

- **ISPs and ASPs should disclose information about their user- or application- based network management and congestion management practices for Internet services in a manner that is readily accessible to the general public**. This information should be made available on network operators' public web sites and through other typically used communications and channels, including mobile apps, contract language, or email. ISPs and ASPs may choose to use a layered notice approach, using a simple, concise disclosure that includes key details of interest to consumers complemented by a more thorough and detailed disclosure for use by more sophisticated users, application developers, and other interested parties. The detailed disclosure should include: descriptions of the practices; the purposes served by the practices; the types of traffic subject to the practices; the practices' likely effects on end users' experiences; the triggers that activate the use of the practices; the approximate times at which the practices are used; and which subset of users may be affected. The disclosures should also include the predictable impact, if any, of a user's other subscribed network services on the performance and capacity of that user's broadband Internet access services during times of congestion, where applicable.

- **Network operators should use accepted industry "Best Practices," standardized practices, or seek industry review of practices.** Network standards setting organizations and technical industry bodies produce considered recommendations of Best Practices and standard practices for a variety of operational issues including congestion and congestion management. Where network operators see the need for an innovative solution that has not been standardized or documented as a Best Practice, these network operators should consider bringing their unique network or congestion management practices to such groups for discussion and documentation.

- **When engaging in a congestion management practice that could have a detrimental impact on the traffic of certain users or certain applications, the practice should be designed to minimize that impact.** Some congestion management practices may cause certain users or certain applications to experience performance degradation. ISPs and ASPs should seek to minimize such degradation to the extent possible while still managing the effects of the congestion that originally triggered the use of the practice.

- **If application-based congestion management practices are used, those based on a user's expressed preferences are preferred over those that are not.** User- and application- agnostic congestion management practices are useful in a wide variety of situations, and may be sufficient to accommodate the congestion management needs of network operators in the

majority of situations. However, at times network operators may choose to use application-based congestion management practices, in which case those that prioritize application traffic according to a user's expressed preferences are preferred over those that do not.

- **If application-based criteria are used by a network operator, they should be tested prior to deployment and on an ongoing basis.** Application-based classification by network operators (e.g., using deep packet inspection) can sometimes be erroneous. If network operators choose to use application-based criteria for congestion management, the accuracy of the classifier should be tested before deployment.

- **ASPs and CDNs should implement efficient and adaptive network resource management practices.** ASPs and CDNs should match use of network resources to the performance requirements of the application. Applications should be designed to efficiently and adaptively use network resources, to the extent feasible given the application's requirements.

F. Port Blocking (Aug 2013) – *Executive Summary and Recommendations*

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.

G. SNMP Reflected Amplification DDoS Attack Mitigation (Aug. 2012) – *Executive Summary and Recommendations.*

A BITAG member and Internet Service Provider (ISP), Comcast, has observed large-scale Simple Network Management Protocol (SNMP) Reflected Amplification Distributed Denial of Service (DDoS) attacks. These attacks are significant and have been observed to result in tens of gigabits to over one hundred gigabits per second of SNMP traffic sent to attack targets from multiple broadband networks. These attacks have been hours long in duration, disruptive for attack targets, and very challenging for targets to mitigate. The conditions that make this attack possible exist on many types networks, regardless of access network technology (DOCSIS, DSL, fiber, etc.), and regardless of geographic location.

The general conditions making this possible include:

- Some networks do not perform ingress filtering, which makes it possible for users of those networks to spoof packets, making it appear that the packets originated elsewhere.
- Networks have hosts that are infected with malware, and are under the control of bot networks.
- Some home gateway devices (a.k.a. routers) ship with SNMP turned on by default, using a well-known community string such as "public."

To conduct the attack, the following steps are taken by an attacker:

- Initiation: An attacker sends instructions to a bot network to conduct the attack. These instructions include the bots to use to distribute the attack, the home gateways to reflect and amplify the attack, and the IP address of the attack target.
- Distribution: Infected hosts participating in a bot network, which happen to be located in a network that cannot or has not taken sufficient steps to prevent spoofing, receive the attack instructions. Thus, one attacker distributes the attack activity to many individual hosts. Each of the multitude of bots sends a small SNMP query to home gateway devices that are listening for particular SNMP queries on their public Internet network interface. This query is forged to make it appear that it was sent from the victim's IP address, so that all responses will be directed to the target rather than back to the bot network's hosts.

- Reflection: Home gateways that were listening for SNMP queries, receive the forged queries from the bot network's hosts. They then send an SNMP response to the target.
- Amplification: The size, in bytes, of the SNMP response is larger than the SNMP query sent by the bot network. So the bot network is able to amplify the amount of data directed at the attack target, compared to a smaller amount of data sent by the bot network.

Device makers as well as Internet Service Providers (ISPs) and Application Service Providers (ASPs) should be aware of this issue and may need to consider a range of potential network management or other responses. The recommendations of the BITAG include:

- End-user devices should not be configured with SNMP on by default.
- End-user devices should not be routinely configured with the "public" SNMP community string.
- ISPs, ASPs, and other network or systems administrators should not routinely use the "public" SNMP community string on an unsecured basis.
- Users should be allowed and encouraged to disable SNMP.
- ISPs should take reasonable steps to prevent address spoofing.
- ISPs may implement appropriately targeted filtering/blocking of SNMP traffic.
- ISPs should be transparent with respect to network management policies that may impact SNMP traffic.
- ISPs should provide mechanisms to re-enable SNMP on a case-by-case basis.
- ISPs and attack targets should be willing to share relevant and non-proprietary information related to SNMP-based attacks with appropriate communities.

H. IPv6 Large Scale Network Address Translation (NAT) (March 2012) – ***Executive Summary & Recommendations***

The Internet is running out of addresses in the format in which they were originally standardized, known as IPv4, due to aspects of that format which constrain the address space to a relatively "small" number of unique addresses as compared to the burgeoning number of devices requiring those same addresses to function on the Internet. A successor address format, IPv6, has been developed to support as many devices as can conceivably be connected to the Internet for the foreseeable future. While the global transition to IPv6 is in progress, it is going to take a number of years to upgrade *all* Internet applications and services, consumer electronics

devices, and networks to support IPv6. The transition period is also likely to be lengthy given that, among other things, IPv4-only equipment is still being manufactured and sold to consumers. As network operators deploy IPv6 technology into their existing IPv4 networks, IPv4 and IPv6 will thus need to co-exist until the demand for IPv4 services diminishes.

Given the amount of time it may take to migrate out of a pure IPv4, or mixed IPv4 and IPv6 network environment, to pure IPv6 service, network operators are employing a variety of techniques to extend the life of IPv4 addressing. One such technique is the use of Large Scale Network Address Translation (also known as Large Scale NAT or LSN). LSN equipment allows a large number of IPv4-enabled end devices to share a single public IPv4 address. Network Address Translation (NAT) functionality has long existed in local/private networks to help network operators manage their network addresses using private address space. NAT functionality is known to adversely impact some Internet applications; wider use of NAT as part of LSN therefore deserves careful examination.

The BITAG is interested in LSN given that IP address sharing is a key tool for extending the life of IPv4 during the transition to IPv6. LSN is likely to affect many players in the Internet ecosystem: ISPs, end users, application providers, equipment vendors, content delivery networks, and third parties such as law enforcement agencies. A broad understanding of problems that may arise has the following benefits: (1) it will help stakeholders to prepare for actions that minimize the impact on end users and applications; (2) it will inform policymakers and regulators of the motivations and trade-offs for the deployment of this technology; (3) it will accelerate the transition to IPv6; and (4) it will more generally help to reduce or preclude friction and/or conflict surrounding use of this technique among stakeholders, as some have argued that Large Scale NAT could be abused by parties for anti-competitive, discriminatory, or other non-technical purposes.

LSN Deployments and Impacts

LSN will be deployed in different ways depending on which IPv6 transition technologies are in use. These alternatives are discussed in the body of this paper. For all of these alternatives, there are a variety of technical implications of LSN for end users, ISPs, and application providers to consider.

The address sharing enabled by LSN use impacts end users in three primary ways: (1) the number of connections available per user is affected, (2) the ability to uniquely identify an end user device solely via the source IP address is lost, and (3) it becomes much more difficult to reach and maintain connectivity to end user devices. All of these impacts are present in local/private network implementation of NATs. Introduction of LSN increases the probability that users will be affected due to sharing of the port number space. The number of users affected by the limitation in port availability may also increase for the same reason. However, note

that a user checking email or performing simple web browsing functions will not be affected by the LSN.

Internet Service Providers (ISPs) electing to use LSNs must balance the impacts of new network infrastructure (operational and capital costs) and of engineering this new infrastructure for scalability, resiliency, security, and capacity, as well as meeting mandates to be able to log individual customer IP address assignments, with maintaining an appropriate level of customer service. In the mobile environment, where every device must be assigned at least one IP address and where simple devices may have limited access to Internet applications, mobile operators have already implemented LSN and faced some of these challenges. However, the continued swift growth in the number of mobile customers and their likely evolution toward expecting wireless Internet service to behave in a manner comparable to wireline service presents new challenges.

LSN can have a wide variety of impacts on applications. These may relate to capacity constraints if the LSN is undersized, the handling of multiple connections to the same application server, the loss of IP-based geolocation capability, new logging requirements, and a variety of other factors.

Recommendations

BITAG has compiled the following recommendations regarding steps that can be taken to help ensure optimal user experience, balanced with efficient LSN deployments and operations:

- **Commit to rapid deployment of IPv6.** The best way to mitigate the impacts of LSN is to reach a state where IPv6 is the dominant addressing scheme. BITAG suggests that ISPs deploy IPv6, that equipment manufacturers support IPv6 in their devices, and that applications sensitive to NAT be supported via IPv6 as soon as possible.

- **Address application impacts of LSN.** BITAG suggests that vendors of LSN equipment adhere to existing requirements [Common requirements for Carrier Grade NAT (CGN)] intended to increase the likelihood that applications will function properly in the presence of LSN. BITAG also suggests that ISPs test their LSN implementations and mitigate application issues prior to deployment, and that application developers use LSN work-arounds or avoid deploying services that do not function properly in the presence of NAT or LSN.

- **Disclose LSN deployment.** To assist with end user troubleshooting, BITAG suggests that ISPs be transparent with respect to the locations and timing of LSN deployment.

- **Provide mechanisms to facilitate LSN traversal to end users.** BITAG suggests that, where feasible, ISPs and equipment vendors support mechanisms to facilitate NAT traversal, including mechanisms for the manual or automatic creation and management of port forwarding rules. Such mechanisms increase the likelihood that applications requiring inbound connections to end users can function across LSN.

- **Provide contact information.** BITAG suggests that ISPs provide a means for application providers to contact them to discuss LSN impacts and possible mitigations.

- **Consider Logging Impacts of Port Allocation.** BITAG suggests that ISPs deploying LSN consider logging and operational impacts when deciding whether to implement a deterministic or dynamic mechanism (or a hybrid of the two) for assigning ports to subscriber sessions.

- **Include Port Number When Logging Activity**. BITAG suggests that Application Providers that maintain a log of user activity include both the IPv4 address and port number in the log. This would ensure that logs accurately reflect the actions of a single ISP customer when IPv4 traffic goes across a LSN.

I. IPv6 AAAA DNS Whitelisting (Sept. 2011) – *Executive Summary and Recommendations*

Internet Protocol version 6 (IPv6) is the new form of addressing for the Internet. IPv6 is being deployed as a result of the exhaustion of the supply of the older form of addressing, Internet Protocol version 4 (IPv4). The Internet is now beginning a global transition from IPv4 to IPv6, although most websites and networks will run both address families in parallel for many years.

This document describes the emerging practice of Internet web sites (or, more broadly, "Internet domains"), such as example.com, selectively returning IPv6-related resources from Domain Name System (DNS) servers. As a result, the domain can enable or disable a network (and, as a consequence, that network's users) from accessing the domain's content over IPv6. This practice is known as "DNS Whitelisting" and is intended as a means to smooth the global transition from IPv4 to IPv6.

DNS Whitelisting was first used by major web content sites. These web site operators observed that when they added IPv6 records to their DNS servers in order to support IPv6 access to their content, a fraction of end users had slow or

otherwise impaired access to this content **[IPv6 Dual-Stack Client Loss in Norway]**.

Major domains are motivated by a desire to maintain a high-quality user experience for all of their users, as well as to shift traffic to IPv6 in a controlled manner. Thus, domains engaging in DNS Whitelisting are attempting to shield users with impaired access from the symptoms of those impairments until they can be remedied. Those domains are also interested in having a degree of control over their own migration to IPv6 that they would otherwise lack without DNS Whitelisting, so that IPv6 traffic can be added gradually as IPv6 operations and network practices mature.

At the same time, critics of the practice have articulated a range of technical and non-technical concerns and are focused on ensuring that DNS Whitelisting, when implemented, is done in a manner that is transparent, non-discriminatory, and not anti-competitive.

Generally speaking, a domain that implements DNS Whitelisting does so manually. This means that the domain manually maintains a list of networks that are permitted to receive IPv6 records (via their DNS resolver IP addresses) and that these networks typically submit applications in order to be added to the DNS Whitelist.

Domain operators foresee that a second phase of DNS Whitelisting may emerge in the future, possibly in the near future. In this new phase a domain would return IPv6 and/or IPv4 records dynamically based on automatically detected technical capabilities, location, or other factors. It would then function much like (or as part of) global server load balancing, a common suite of practices already in use today. Furthermore, in this second phase, networks might be added to and removed from a DNS Whitelist automatically, and possibly on a near-real-time basis. This means that networks may no longer need to apply to be added to a whitelist, which could alleviate some of the issues addressed herein.

Since this future phase has yet to emerge, this document and its suggestions apply only to the DNS Whitelisting practices currently known to the BITAG. The BITAG reserves its opinion on future alternative practices until they can be articulated and evaluated.

Although current implementations are not perceived to have this impact, the BITAG is interested in this issue on that grounds that, without careful and monitored deployment, some whitelisting services could in the future be viewed as anti-competitive, discriminatory or in violation of some other public policy objective. The practice may be viewed as controversial and the manner in which it is employed could result in concerns or complaints. As a result, it is important to inform the public and policymakers about why DNS Whitelisting is

used and how it functions, to identify concerns surrounding its use, and to outline some potential implementation steps that domains could take to minimize the risk of complaints and controversy.

The BITAG has formulated a set of suggested practices regarding the implementation of DNS Whitelisting to help reduce such complaints, acknowledging that practices may vary as a domain moves from experimentation with whitelisting and IPv6 to a point of operational stability. The suggested (and voluntary) practices are as follows:

- **Limit the Duration of the Use of DNS Whitelisting** – The BITAG suggets that domain operators use DNS Whitelisting as briefly as possible for those critical domains that will encounter significant end user IPv6-related impairments. The BITAG recognizes that the primary purpose for DNS Whitelisting is to permit IPv6 address service to networks that are essentially problem-free while preferring IPv4 for those that are not and that the duration of Whitelisting is tied to that assessment.

- **Transparently Publish Policies and Processes** – The BITAG suggests that domains make their policies and processes easy to find and understand. The contact persons should be easy to find and reach as well.

- **Clearly Describe Decision-Making Criteria –** The BITAG suggests that policies and procedures that a domain publishes include criteria used to make DNS Whitelisting and de-whitelisting decisions.

- **Use Primarily Quantitative Decision-Making Criteria –** The BITAG suggests that domains use quantitative data (such as the number of IPv6-related impairments or the performance of IPv6 network routes) to make whitelisting and de-whitelisting decisions.

- **Set and Publish Service Level Goals for Decision-Making** – The BITAG suggests that domains publish clearly detailed and timely service level goals for how long the DNS Whitelisting decision-making process will take.

- **Specify an Appeals Process** – The BITAG suggests that domains may wish to consider specifying a process for networks to appeal both whitelisting and de-whitelisting decisions.

- **Maintain Updated Contacts for Whitelisted Domains** – The BITAG suggests that domains establish a list of contacts for whitelisted organizations and adopt practices that assure that the list is current.

- **Set a Joint-Troubleshooting Interval Before De-Whitelisting Occurs** – The BITAG suggests that domains set a reasonable period of time and process for a

whitelisted party to resolve any problems that may arise that could lead to de-whitelisting. The BITAG recognizes that there may be emergency conditions that require immediate action.

- **Transparently Publish Whitelisted Parties** – The BITAG suggests that domains identify the networks that are currently listed in their DNS whitelists.

- **Openly Share IPv6-Related Impairment Statistics** – The BITAG suggests that domains share detailed statistics about IPv6 impairments with any party (campus network, enterprise, ISP, etc.) that may be affected by DNS Whitelisting. The BITAG recognizes that privacy concerns may limit the kinds of data that can be shared.

- **Detect and Notify End Users with IPv6-Related Impairments** – The BITAG suggests that, where practical, domains take reasonable steps to detect IPv6-related impairments and take reasonable steps to communicate this in an easy-to-understand way to affected users.