[Comments due on 18 July 2016 at **5 AM** Pacific.
Comments are to be submitted by email to NN-Consultation@berec.europa.eu with messages not to exceed 2MB]

To: Dr. Wilhelm Eschweiler, Chairman, Body of European Regulators for Electronic Communications
From: Dr. Paul Vixie, CEO and Chairman, Farsight Security, Inc.
Date: 15 July 2016
Subject: Comments on draft BEREC Guidelines on implementation of net neutrality rules

Dear Dr. Eschweiler:

Please consider the following comments from Farsight Security, Inc., as input to the draft BEREC Guidelines relating to the implementation of net neutrality rules, per the public announcement at http://berec.europa.eu/eng/news_consultations/ongoing_public_consultations/3771-public-consultation-on-draft-berec-guidelines-on-implementation-of-net-neutrality-rules

I. CONTEXT FOR THESE COMMENTS

(1) My background:

I am Chief Executive Officer and Chairman of the Board of Farsight Security, Inc. I've previously served as President, Chairman, and Founder of Internet Systems Consortium (ISC), as President of MAPS, PAIX and other businesses, as CTO of Abovenet/MFN, and serve on the boards of several for-profit and non-profit companies. I have previously served on the ARIN Board of Trustees, including serving as Chairman in 2008 and 2009, and I am a founding member of ICANN Root Server System Advisory Committee (RSSAC) and ICANN Security and Stability Advisory Committee (SSAC). I operated the ISC's F-Root name server for many years, and I am a member of Cogent's C-Root team. I'm also a sysadmin for a leading industry cybersecurity information sharing forum, OpSec Trust.

I've been contributing to Internet protocols and UNIX systems as a protocol designer and software architect since 1980. I wrote Cron (for BSD and Linux), and am considered the primary author and technical architect of BIND 4.9 and BIND 8, and I hired many of the people who wrote BIND 9. I've authored or co-authored a dozen or so RFCs, mostly on DNS and related topics, and wrote Sendmail: Theory and Practice (Digital Press, 1994). My technical contributions include DNS Response Rate Limiting (RRL), DNS Response Policy Zones (RPZ), and Network Telemetry Capture (NCAP). I earned my Ph.D. from Keio University for work related to DNS and DNSSEC, and I was named to the Internet Hall of Fame in 2014.

This broad technical- and Internet governance-related background, and my roles leading innovative and successful Internet technical companies, gives me an expert's perspective from which to review and comment on BEREC's draft network neutrality rules. The remarks below are offered in my capacity as Farsight CEO and Chairman of the Board, and reflect both my own personal perspective on these matters and Farsight Security, Inc.'s official company perspective.

(2) Some Background Information About Farsight Security, Inc.:

Leveraging our deep Domain Name Systems (DNS) expertise, Farsight Security, Inc., offers real-time Passive DNS solutions that provide critical context to significantly increase the value of prepackaged reputation & threat feeds, and other threat intelligence. The availability of timely and relevant security-related data is the key to establishing tactical superiority in any cyber engagement. The entire Farsight Security organization is focused on increasing the availability, variety, volume, quality, breadth, and relevance of the network telemetry data we deliver. Our coordinated efforts allow our customers to increase the variety and effectiveness of their network protections and countermeasures, which can now often even be deployed before attacks are initiated against them. At Farsight Security, we are committed to finding new ways to secure the world's digital infrastructure *while fully respecting and protecting the privacy of all law-abiding Internet users.* More information about Farsight Security, Inc. can be found online at https://www.farsightsecurity.com

(3) Our Interest In The European Union's Work On Network Neutrality:

Since Farsight is a San Mateo, California, U.S.A.-based company, we believe it is important to explain why we are weighing-in on a European Union network neutrality rulemaking proceeding.

While Farsight is based in California, we have numerous international customers, as well as international sensor operator partners and international contractors. We're an international company. Moreover, Internet traffic pays little heed to political geography: a web page hosted by an ISP in Paris, France, ISP is as readily accessed as a web page at an ISP in Pomona, California. Therefore, these proposed rules, while of, for, and by the EU, nonetheless directly affect us, and our company's products, services, and operations.

We also recognize the EU's leadership, and the reality that policies pioneered in and by the EU often ends up becoming model policies for the rest of the world. This leadership and influence further multiplies our interest in your network neutrality work.

Substantively, and if we may presume to speak plainly, our biggest concern is that the EU's proposed network neutrality rules, while obviously carefully prepared and very well-intentioned, may be misinterpreted or badly implemented through no fault of your own. If that were to occur, critical network operations might be hindered, and the safety, security, stability, and interoperability of the Internet might be negatively and even potentially irrevocably impacted.

Given the above rationale, we hope that you will consider our input when making final clarifications to your current draft rules. We understand that our comments will be made publicly available, and you have our permission to use our feedback as may best meet BEREC's rulemaking needs.

[Editor's Note: Our comments below are keyed to specific paragraphs as they appeared in the report that's available online at http://berec.europa.eu/eng/document_register/subject_matter/berec/download/0/6075-draft-berec-guidelines-on-implementation_0.pdf]

II. PRESERVATION OF OPERATIONALLY-CRITICAL NETWORK SECURITY-RELATED CAPABILITIES

(1) We Must Preserve ISPs' Ability to Filter Spam, Malware, Phishing and Other Unwanted Network Traffic

Farsight supports non-discriminatory handling of all **lawful** and **wanted** network traffic.

However, having said that, we are equally emphatic that network neutrality rules must not inadvertently or accidentally erode our ability to block spam, malware, phishing, and other **illegal** or **unwanted** network traffic.

Specifically, if read in isolation and without caveats, paragraph 49 from the consultation document would appear to prohibit any filtering of network traffic whatsoever, even if that traffic is illegal or unwanted:

> *[National Regulatory Authorities] should ensure that traffic on an IAS is managed:*
>
> - *"without discrimination, restriction or interference";*
> - *"irrespective of the sender and receiver, the content accessed or distributed, the appl*
>
> *or*
>
> *services used or provided, or the terminal equipment used".*

Fortunately, we note that Article 3(3) letter (b) and paragraph 79-81 provides an exemption for a variety of cybersecurity practices. Our concern is that this section merely does so in passing, and just by way of an incomplete list of examples. This may result in potential ambiguity as to the permissible techniques and the permissible extent of such protective measures. We would prefer to see BEREC provide greater specificity and clarity in this area.

For example, in paragraph 79, four categories of attacks are called out: DDoS attacks, spoofed traffic, "hacking attacks," and distribution of "malicious software, viruses, etc." A number of attacks are conspicuous by their absence, including spam, phishing, port scanning, man-in-the-middle (MITM) attacks, eavesdropping on network traffic and so forth. We urge you to clarify the full range of network and system security threats which may be mitigated without infringing an

ISP's network neutrality obligations, or to at least clarify that the list of network and system security threats enumerated are meant to be illustrative rather than exhaustive.

(2) We Must Preserve ISPs' Flexibility With Respect to HOW They Filter or Otherwise Manage Unwanted Traffic

Similarly, in paragraph 80, a variety of blocking strategies are outlined, including blocking IP addresses or specific port numbers. We draw your attention to Recital 2, which states that

> *"The measures provided for in this Regulation respect the principle of technological neutrality, that is to say they neither impose nor discriminate in favour of the use of a particular type of technology."*

We urge adherence to that principle when it comes to the technical options that are available for ISPs use in managing cyber security threats.

Thus, while my friend and colleague Eric Ziegast and I are credited with creating the first Internet block lists (see https://en.wikipedia.org/wiki/DNSBL), and it is certainly true that blocking IP addresses or blocking specific ports remains a popular approach today, blocking IP addresses or port numbers is NOT the only potential approach to cybersecurity threats today. Blocking IP addresses and blocking port numbers may not even be the best approach any longer. For example:

-- Some spam is most readily blocked by using domain-based block lists such as the Spamhaus DBL or SURBL

-- Other unwanted traffic may be most readily identified and blocked based on distributed checksum approaches, as pioneered by myself and my friend and colleague Vernon Schryver (see https://en.wikipedia.org/wiki/Distributed_Checksum_Clearinghouse)

-- Still other unwanted content (such as some types of malware) might be identified based on signatures associated with executables, or the hash of the executables (see http://www.team-cymru.org/MHR.html)

-- Some ISPs employ DNS firewalls which leverage DNS Response Policy Zones (see https://dnsrpz.info/)

-- Another approach employs a host of different rules, some negative and others positive, which in aggregate determine the spamminess of a particular message (see http://spamassassin.apache.org/tests_3_3_x.html)

That short list is far from an exhaustive enumeration, but it does serve to illustrate the point that there are many different ways to tackle security threats confronting ISPs and their customers. BEREC should refrain from implicitly picking "winners" and "losers" from among all potential approaches, particularly since the needs and approaches the cybersecurity employs continually evolve. Allow ISPs to choose the approach that works for them and their circumstances.

(3) ISPs Must Remain Able to Instrument Their Networks So As To Be Able To Detect and Respond To Threats

In paragraph 81, you discuss the important role of netwok monitoring when it comes to tackling cybersecurity threats. We applaud this insight. A well-instrumented network is indeed the first step toward securing any production network. Without appropriate monitoring, you're effectively hurtling down the road with a blacked-out windshield, no headlights and a missing rearview mirror, unaware of either what's ahead or what you've just passed.

We'd only ask that you clarify that a **broad range of network monitoring approaches may be fruitful,** and a **broad range of network monitoring approaches are permissible** for ISPs to use under EU current and contemplated regulations, at least provided that they're used by an ISP's security team to identify vulnerabilities in their own systems or networks, and/or to keep ISP facilities and customers safe.

A non-exhaustive list of some types of network and system monitoring approaches that some ISPs have found helpful include:

-- Network flow collection and analysis (e.g., Netflow, Jflow, Sflow, etc.)
-- Intrusion Detection Systems (IDS) reporting and alerting (Snort, Bro, Suricata, etc.)
-- Active scanning tools (such as Metasploit, Nessus, Nikto, etc.)
-- Centralized syslogging and analysis (popular tools include Splunk; Elasticsearch, Logstash and Kibana; etc.)
-- Logging of DNS cache miss traffic collected above recursive resolvers,
-- SNMP monitoring for things like traffic levels in octets and packets, network errors, etc., and
-- Antimalware scanner package reporting.

(4) The Specific Issue of Controlling Spoofed Network Traffic ("Source Address Validation")

As we're sure you're aware, a huge problem facing ISPs, enterprise networks, and even national governments is **Distributed Denial of Service (DDoS) attacks.** While these attacks can take a variety of forms, ranging from simple brute force volumetric packet flood to more subtle state exhaustion attacks (such as "half-open" SYN floods), it is critical that ISPs have the ability to protect themselves -- and the ability to help protect the rest of the Internet.

For example, consider DNS amplification attacks. In this type of DDoS attack:

-- An attacker forges DNS query traffic so that it misleadingly appears to come from the attacker's intended victim
-- That forged traffic is then directed to misconfigured DNS servers ("open recursive resolvers") located all around the Internet.
-- When those forged queries are received by the open recursive resolvers, they reply with a (typically large) response to those questions -- even though the victim never asked those queries. This response traffic, in aggregate, can total gigabits per second and can render the victim's network totally unusable.

Firewalls and router ACLs are too late to help: by the time the DNS amplification traffic hits those filters, chokepoint network links have already been totally saturated.

ISPs play a potentially critical role when it comes to stopping these sort of attacks: if ISPs do Source Address Validation (SAV, see https://tools.ietf.org/html/rfc6959), they can ensure that forged traffic never leaves their network, thereby helping to stop these type of attacks in their tracks.

We notice that DDoS attacks and blocking spoofed traffic are both mentioned in paragraphs 79-80, Bravo! We would be reassured, however, if the proposed rules clarified explicitly that ISPs are allowed, nay, **urged** to prevent spoofed traffic from leaving their networks. This is an utterly rock-bottom basic network management issues, totally non-discriminatory, and critical to helping to eliminate the scourge of Distributed Denial of Service attacks. Without this clarification, we're concerned that some sites may not be willing to begin doing Source Address Validation, and some sites that are currently doing SAV may stop using it, fearful of running afoul of EU network neutrality requirements.

(5) Online Child Sexual Abuse Materials and Online Illegal Hate Speech Materials

The other areas where we hope you will be utterly unambiguous when it comes to an ISP's responsibilities are:

-- Online Child Sexual Abuse Materials (sometimes mistakenly referred to as "child pornography"), and

-- Online illegal hate speech materials (most characteristically consisting of incitement to perform acts of terrorist violence, but extending to additional content as well, as discussed in "European Commission and IT Companies announce Code of Conduct on illegal online hate speech," see http://europa.eu/rapid/press-release_IP-16-1937_en.htm , 31 May 2016)

We believe that any ISP action taken to combat these two scourges is consistent with Article 3(3) letter (a) and Recital 13, but urge you to eliminate any potential ambiguity or lack of clarity with respect to these two specific areas of special concern.

ISPs must have no doubts about their ability to act to tackle these two urgent areas.

III. REDUCING CONSUMER CHOICES?

While we support non-discriminatory treatment of all lawful and wanted network traffic, we also ardently support a diverse range of consumer options. We were thus pleased to note Article 3(2), which provided:

> *Agreements between providers of internet access services and end-users on commercial and technical conditions and the characteristics of internet access services such as price, data volumes or speed, and any commercial practices conducted by providers of internet access services, shall not limit the exercise of the rights of end-users laid down in paragraph 1.*

Put simply, that provision appears to say (at least as we read it), "Consumers should be able to choose the Internet service offering they want." We think that's as it should be. Some people may want to fly first class, while others may be fine flying coach. Some are willing to spend the money to drive a brand new sports car, while others may want a used commuter vehicle, spending the difference elsewhere.

Sadly, while the proposed regulations are generally consistent with Article 3(2), things go "off the rails" with the discussion of "zero rating" at paragraph 37-39. Those paragraphs state:

> *37. There is a specific commercial practice called zero-rating. This is where an ISP applies a price of zero to the data traffic associated with a particular application or category of applications (and the data does not count towards any data cap in place on the IAS).There are different types of zero-rating practices which could have different effects on end-users and the open internet, and hence on the end-user rights protected under the Regulation.*
>
> *38. A zero-rating offer where all applications are blocked (or slowed down) once the data cap is reached except for the zero-rated application(s) would infringe Article 3(3) first (and third) subparagraph (see paragraph 52).*
>
> *39. The ISP could either apply or offer zero-rating to an entire category of applications (e.g. all video or all music streaming applications) or only to certain applications thereof (e.g. its own services, one specific social media application, the most popular video or music applications). In the latter case, an end-user is not prevented from using other music applications. However, the zero price applied to the data traffic of the zero-rated music application (and the fact that the data traffic of the zero-rated music application does not count towards any data cap in place on the IAS) creates an economic incentive to use that music application instead of competing ones. The effects of such a practice applied to a specific application are more likely to "undermine the essence of the end-users' rights" or lead to circumstances where "end-users' choice is materially reduced in practice" (Recital 7) than when it is applied to an entire category of applications.*

In this case, we believe BEREC's reasoning is flawed. There are many examples where one product or service is provided at no (or virtually no) out-of-pocket cost, yet consumers still elect to pay a premium for a product they like better. Some illustrative examples from other sectors:

-- Broadcast music (via over-the-air FM radio) is free once you've purchased a radio, yet Apple has nonetheless managed to sell well over 25 billion songs to users via iTunes. [see https://techcrunch.com/2013/02/06/charting-the-itunes-stores-path-to-25-billion-songs-sold-40-billion-apps-downloaded-and-beyond/]

-- Similarly, while potable tap water is available from the tap throughout virtually all of the EU at negligible incremental cost, EU consumers still choose to purchase many many liters of bottled water per capita per annum. [See http://www.efbw.eu/index.php?id=90]

If we were to extend BEREC's network reasoning to the music or water sectors, consumers would have to be "protected" from the "economically-corrosive" effects of free over-the-air music and (virtually-)free tap water, presumably to ensure that Apple and Evian enjoy an "economically level playing field." The prospect that these and similar firms might be economically disadvantaged--or that ISPs would be disadvantaged by typical "zero-rated" offerings--is ludicrous. Don't deny users the opportunity to enjoy a bargain if one is available. Forbidding minor bundled offerings of this sort actually directly **reduces** consumer choice, while perversely claiming the expansion of consumer choice as a justification!

As a thought experiment, let us turn things on their head. Imagine a standardized portfolio of ISP services required to be offered by each approved provider. Let it be required that those services be offered throughout the whole of the EU, and at a uniform ("postalized") price regardless of where they may be purchased. Each provider would be required to meet the *minimum* characteristics associated with each tier of service, but would be free to offer additional services -- what we might call zero-rated services -- in order to compete for consumer market share. Is this somehow economically unfair? We think not. Consumers should be free to choose the *combination* of basic service, price, *or special features* that best meet their needs, and providers should be free to offer unique combinations of features that the provider believes consumers will find most attractive. ISP service is, *or should be allowed to be,* more than just a matter or price and basic bandwidth.

IV. LIMITING MANAGEMENT OF NETWORK CONGESTION -- AND INCREASING CONSUMER COSTS?

Paragraph 89 states that:

> [National Regulatory Authorities] should monitor that ISPs properly dimension their network, and take into account the following:
>
> • if there is recurrent and more long-lasting network congestion in an ISP's network, the ISP cannot invoke the exception of congestion management (ref. Recital 15);
>
> • application-specific congestion management should not be applied or accepted as a substitute for more structural solutions, such as expansion of network capacity.

Let us contrast two different models for building out a network.

In the first model, we assume that each customer might routinely and continually uses 100% of their provisioned edge capacity for whatever purpose they desire. In that scenario, upstream links need to be provisioned to be at least as large as the sum of all subordinate links. For example, an ISP with 100 customers, each connected at 100 Mbps, would need to provision a 10Gbps uplink to ensure that they can always accommodate the maximum possible aggregate customer-offered load. Naturally, most of the time, a substantial fraction of that shared 10Gbps link will go unutilized as a matter of factual reality, provisioned "unnecessarily" when viewed against actual observed traffic levels.

In the second model, taking advantage of statistical multiplexing, packet-switched networks leverage the fact that most customers do not routinely and fully utilize their link's potential capacity. Some customers may talk now for a bit, and then go silent, while other customers may talk a bit later, and then go silent in turn, etc. As modeled, on average, some fraction of the maximum potential capacity is all that will *normally* be needed to service the offered demand. However, if a "roll of the dice" happens to be extreme, there may be times when the shared upstream capacity is NOT sufficient, and congestion will result.

Obviously, most consumer connections employ the second model rather than the first. So why use this second model? The economies associated with the second model get reflected both in what consumers pay, and in what consumers are allowed to do: a statistical model often is configured to allow users to potentially "burst" above their nominal capacity.

But now, the requirements of paragraph 89, as quoted above, requires solely that ISPs respond to persistent congestion by adding capacity (e.g., moving from the second model to something closer to the first model), even if the need for such upgrades is driven largely by the behavior of just a handful of customers (perhaps because they're operating a server or doing peer-to-peer file sharing).

Adding capacity has associated costs. So who will pay? Those few whose disproportionate usage is causing the need for that premature upgrade? No. ALL users will be made to suffer (by paying more) to accommodate the disproportionate bandwidth appetites of the few. ALL users will also typically have lower flexibility, e.g., less ability to occasionally burst above strictly committed bandwidth limits. In our opinion, both of these outcomes are unfortunate and totally avoidable if reasonable usage limitations could once again be allowed. "Trimming the tail" of the usage distribution by curtailing the extreme actions of a few is far more cost effective -- and far more fair -- than penalizing all users for the "sins of a few."

V. CAPACITY ISN'T THE ONLY THING NEEDED TO ACTUALLY REALIZE WIRE-RATE THROUGHPUT

Much of the proposal under discussion is focused on ensuring that some traffic isn't "slowed down" (or otherwise disadvantaged) relative to some other comparative traffic.

This focus is laudatory, but neglects an important reality: most systems are poorly tuned and totally unprepared to take full advantage of high capacity connections. That is, even if you're connected by a lightly-loaded 100Mbps or 1000Mbps (Gbps) connection, that **doesn't** mean that you'll see wire-rate throughput. Most hosts simply aren't optimized to deliver wire-rate high-bandwidth throughput for reasons described at https://fasterdata.es.net/host-tuning/

We believe that if the goal is to help end users actually achieve full wire-rate throughput on high capacity connections, ISPs will need to help customers tune their systems. We do not believe that this is an aspect of the problem that's currently getting the attention it deserves, and this is not a problem that will somehow spontaneously "self-remediate." If it is left unaddressed, many users may have nominally "fast" connections, but only a rare few will achieve wire-rate Internet throughput, best efforts of policy makers notwithstanding. Policies, without tuning assistance, will not eliminate throughput inequalities.

We also believe that routine wide-area throughput measurements using iperf3 (see https://iperf.fr/) are critical to empirically demonstrating what wide-area throughput can be achieved ("show me, don't just tell me"), and provides critical early warnings of problems if problems do arise. Throughput measurements are an integral part of the measurement programs supported by most academic high performance networks and specialized agency mission networks, such as ESNet. (See https://fasterdata.es.net/performance-testing/perfsonar/)

VI. PERSONAL DATA

Article 3(4) provides that

> *Any traffic management measure may entail processing of personal data only if such processing is necessary and proportionate to achieve the objectives set out in paragraph 3. Such processing shall be carried out in accordance with Directive 95/46/EC of the European Parliament and of the Council. Traffic management measures shall also comply with Directive 2002/58/EC of the European Parliament and of the Council.*

We believe that BEREC should treat individual personal data and aggregated personal data differently. In particular, we recommend that the EU take notice of the U.S. Federal Communication Commission's recent Notice of Proposed Rulemaking relating to "Protecting the Privacy of Customers of Broadband and Other Telecommunications Services," FCC 16-39, http://transition.fcc.gov/Daily_Releases/Daily_Business/2016/db0401/FCC-16-39A1.pdf

In paragraph 153, the Commission describes its perspective on aggregate customer information, stating that:

> *Because of the complexity of the issues surrounding aggregation, de-identification, and re-identification of the data that [Broadband Internet Access Service] providers collect about their customers, we propose to address separately the use of, disclosure of, and access to aggregate customer information. Consistent with reasonable consumer expectations, existing best practices guidance from the FTC and NIST, and Section 222(c)(3)'s treatment of aggregate [Customer Proprietary Network Information], we propose to allow BIAS providers to use, disclose, and permit access to aggregate customer [Personal Information] if the provider (1) determines that the aggregated customer PI is not reasonably linkable to a specific individual or device; [other arguably unnecessary potential requirements omitted here]*

We support an explicit "carve out" for aggregate personal data, particularly for things like DNS cache miss traffic collected from above large shared recursive resolvers. Such data is extraordinarily valuable for fighting cyber crime, and in this aggregated form, is impossible to attribute to particular end-users (all queries simply appear to come from the recursive resolver itself).

Currently, ambiguity about what is and isn't allowed when it comes to sharing aggregated data deters at least some ISPs from participating in fundamental cybersecurity data sharing efforts. We urge BEREC to make it clear that ISPs can confidently share **aggregated** personal data without risk or worry of being subject to sanctions.

VII. CLOSING

Thank you for performing this important rule making, and thank you for the opportunity to comment on these proposed regulations. Please feel free to get in touch by email (vixie@fsi.io) or by phone (1-650-489-7919, U.S. Pacific Time) if we can help in any further way.

Sincerely,

/s/

Dr. Paul Vixie, CEO and Chairman of the Board
Farsight Security, Inc.