

The Internet Protocol Journal

March 2016

Volume 19, Number 1

*A Quarterly Technical Publication for
Internet and Intranet Professionals*

In This Issue

From the Editor	1
What's in a DNS Name?	2
QoS and QoE	14
Fragments	41
Call for Papers	42
Supporters and Sponsors	43

FROM THE EDITOR

This *Internet Engineering Task Force* (IETF) has just completed its meeting in Buenos Aires, Argentina. This meeting was the first time the IETF has met in South America, and while Buenos Aires is “far away” from many parts of the world, the technical community seems to agree that this historic meeting was well worth the journey. You can read more about this meeting on the IETF and Internet Society websites, as well as in the latest issue of the *IETF Journal*, which on this special occasion is available in both English and Spanish.

This year, the IETF will meet in Berlin, Germany, in July and in Seoul, Korea, in November. If you are involved with developing or deploying Internet protocols, I recommend that you attend an IETF meeting if you have not done so already. There are also ways to participate remotely in these meetings if you are unable to attend in person, and the IETF now has many resources for first-time attendees at its meetings. For more information about the IETF, visit <http://ietf.org>

The *Domain Name System* (DNS) is the “human face” of the Internet, allowing us to use terms such as **facebook.com** or **isoc.org** to access a service over the Internet. A small number of these names are “reserved” in the sense that they do not appear in the global DNS system. In our first article, Geoff Huston discusses the *Special-Use Domain Name* registry, which has sparked quite a bit of debate in recent months.

Quality of Service (QoS) has been discussed in several articles in this journal over the years. In our second article, William Stallings and Florence Agboma describe QoS and the more recent concept of *Quality of Experience* (QoE) as it relates to IP networks.

We would like to remind you that this journal depends on the generous support of numerous individuals and organizations. If you would like to help support IPJ, please contact us for further details. Comments, suggestions, book reviews, and articles are always welcome. Send your messages to ipj@protocoljournal.org

—Ole J. Jacobsen, Editor and Publisher
ole@protocoljournal.org

You can download IPJ
back issues and find
subscription information at:
www.protocoljournal.org

ISSN 1944-1134

What's in a DNS Name?

by Geoff Huston, APNIC

What's the difference between `.local` and `.here`? Or between `.onion` and `.apple`? All four of these labels are capable of being represented in the Internet's *Domain Name System* (DNS) as *generic Top Level Domains* (gTLDs), but only two of these are in fact delegated names. The other two, `.local` and `.onion` not only don't exist in the delegated name space, but by virtue of a registration in the *Internet Assigned Numbers Authority* (IANA)'s *Special Use Domain Name* registry^[1], these names cannot exist in the conventional delegated domain name space.

It seems that Internet does not have a single coherent name space, but instead it has a name space that contains a number of silent and unsignalled fracture lines, and instead of being administered by a single administrative body there are numerous people who appear to want to have a hand on the tiller! Let's look at the Internet's domain name space and try to gain some insight as to how we've managed to get ourselves into this somewhat uncomfortable position.

A Very Brief History of the DNS

It is probably an impossible challenge to consider many years of development and take the outcome of many discussions, conferences, as well as countless millions of mail messages and generate a brief but complete history of the domain name system. Here I'll offer a personal interpretation of what I recall, supplemented with reference to numerous useful sources, but nevertheless it's still a somewhat subjective narrative.

A good place to start is probably RFC 920^[2], authored by Jon Postel and Joyce Reynolds, and published in October 1984. The name model of the Internet had broken away from many other contemporary "flat" or limited hierarchy naming models used in other computer networks by adopting a hierarchical name scheme that imposed no a priori limit on the depth of the hierarchy. This meant that the apex level of the name hierarchy could be limited to a number of generic category names, leaving the lower levels of the name space hierarchy to be populated by individual name instances.

This document, RFC 920, specified a division of the apex level of the name space into a small set of so-called *Top Level Domains*. These were the category-based names of `.com`, `.edu`, `.gov`, `.mil` and `.org`, the collection of two-letter country codes as administered by the *International Organization for Standardization* (ISO) and published as ISO-3166^[22], and a temporary name of `.arpa`.

By the time RFC 1034^[3] was published in 1987, there was no distinction drawn between the name space itself and the technology of resolution of these names.

The name space and the name resolution technology that operated on this name space was collectively referred to as the *Domain Name System* (DNS). At the time, the name space was a collection of top-level names overseen by the IANA. Even in those early days there was pressure to expand the set of delegated top-level domains. Initially **.net** was added, then **.int**, but these additions appeared to exacerbate the issue rather than relieve these growing pressures. As well as efforts to clarify the nature and administration of the domain name space at the time^[4], the debate over who and how the name space could be further expanded continued as a sometime vexatious topic, particularly as the set of stakeholders and interested parties began to grow. Subsequent investigation to expand the DNS name space was undertaken by the *International Ad Hoc Committee* (IAHC)^[5], sponsored by the IANA, the *Internet Architecture Board* (IAB) and the Internet Society, with membership drawn from a number of bodies including the *Telecommunication Standardization Sector of the International Telecommunication Union* (ITU-T) and the *World Intellectual Property Organization* (WIPO). This committee produced a report that advocated the limited expansion of the collection of gTLDs by adding a further seven top-level labels to the domain name space, as well as proposing some structural changes in the name registration function to delineate the roles of name registrars and name registry operators.

However, perhaps of greater interest were other activities that were underway at the same time as the committee was undertaking its investigation. In 1995 the *National Science Foundation* (NSF) had authorized a company called Network Solutions to operate the names registry for the Internet, and permitted the company to charge an annual fee to maintain a name registry entry, and to keep the proceeds from this operation. This situation caused a significant level of discontent, as there was a general perception that the registration fee was unrelated to the cost of operation of the registry and that the registry operator was exploiting a de facto monopoly position to its benefit. A number of activities emerged in alternate name systems. These alternate name systems used the same name structure, and the same name resolution tools, but used a different set of “root” name servers. These systems were so defined to sit alongside the incumbent name system, but added a number of additional top-level labels (see, for example the Wikipedia account of AlterNIC’s brief history^[5]). At issue here was the coherence of the Internet’s name system. A user whose domain name resolvers were positioned within the name space as defined by one of these alternate name systems could use a name in a communication to another user where the same name may have been defined in a different name system and resolved in an entirely different manner.

In May 2000 the IAB published RFC 2826^[7], which argued strongly for the presentation of a single root system and thereby argued strongly for a single coherent name system: “There is no getting away from the unique root of the public DNS.”

Rather than having the DNS name space grow from the “bottom up” in several uncoordinated grass roots efforts to expand the name space, and allowing each effort to fail or survive on the level of public interest and commercial uptake, the IAB was espousing a view that any such expansion of the name space was to be a top-down effort. All such new top level names were to be implemented in a coherent manner such that all such names were visible to all Internet users at the same time. Any expansion of the domain name space was intended to be a process that included all parts of the Internet, and that at all times all public DNS names were to be equally and uniformly available to all users.

However, at much the same time as this statement was made, mid-2000, the IAB was also attempting to extricate itself and the *Internet Engineering Task Force* (IETF) from the fraught debate about the accountability of the IANA, and the nature of the role of the US Government agencies that had been funding the work of the IANA. This debate also folded in the discussion of the further expansion of this domain name space. Evidently many people at the time were interested in seeing a distinct community of interest focus on the issue of the policy of the domain name space in a manner similar to the evolution of the addressing community and the emergence of the *Regional Internet Registry* model in the 1990s. In June 2000 the IAB entered into an agreement with The *Internet Corporation for Assigned Names and Numbers* (ICANN) that effectively passed over the administrative purview of the domain name space, apart from “assignments of domain names for technical uses,” to ICANN, RFC 2860^[8].

From that point onward the focal point for the debate about the expansion of the name space, and the related debate about the monopoly position of Network Solutions was essentially ICANN. Over the ensuing years ICANN made a number of decisions in the interest of addressing perceived needs that were voiced from the community of interest. The roles of the registry and the front-end registrar function were cleaved apart and competition between registrars allowed the retail price of name registrations to be subject to competitive market pressures. In addition, a number of new gTLDs were added in a relatively ponderous and deliberative process. In 2000 the gTLDs of **.aero**, **.biz**, **.coop**, **.info**, **.museum**, **.name** and **.pro** were added to the delegated name set of the root zone of the Domain Name System. Four years later a second round saw the addition of **.asia**, **.cat**, **.jobs**, **.mobi**, **.port**, **.tel**, **.travel** and **.xxx**.

This conservative approach to augmenting this root zone delegated name set changed with the so-called “new gTLD” program, that started in 2008 with the adoption by the ICANN Board of a number of policy recommendations relating to the expansion of the gTLD delegated name space, and the subsequent 2011 launch of this program.

The application window opened on January 12, 2012, and ICANN received 1,930 applications for new gTLDs. On December 17, 2012, ICANN held a prioritization draw to determine the order in which applications would be processed during initial evaluation and subsequent phases of the program. One view is that these names were effectively sold into the market, with an application fee of \$185,000 USD per name. An alternate view was that the application process now entailed significant levels of analysis of the impact on the broader environment, including considerations of competition, security, collisions, potential trademark infringement and similar subjects, and that this fee was intended to cover some portion of the costs of this investigation of the potential impact of the delegation of this particular name as a new gTLD.

Name Tensions and Collisions

Expanding the gTLD name space did not address all of the outstanding issues, and to some extent these tensions were exacerbated by the chosen mechanism for this expansion. The new names and their “owners” were defined essentially by the actions of bidding for names. Without putting too fine a point on it, the expansion of the Domain Name System was passed to a market-based mechanism that was based on foundations of a commercial model of monetization of the name space. This shift appears to have prompted other forms of use of non-delegated top-level domain names to be a little more visible.

There are a number of examples of this change in the landscape of the domain name space.

Local Names

The first of these is the use of the name space in *private* domains. Although the public name space is held together with the coordinated set of root name servers and a common convention that all public name resolvers use these root name servers to establish content, this is only a convention for the public name space. Within private environments it is quite common to see name servers that define a local name environment as a local convenience. For example, you could call the local data server in your home network **server.home**. Not only is that name convenient for the home user, it's convenient for a vendor of home equipment, who can preconfigure server equipment and use these local private names in a pre-configured mode. There are a many names that are commonly used in private environments, probably as a result of vendors in this market domain adopting particular name conventions. The names **.home**, **.hometown**, **.belkin**, **.lan**, **.dlink**, and **.local** are all popular names in locally defined private DNS domains^[9].

What happens if ICANN were to delegate a new gTLD that was the same as a name that enjoyed considerable levels of private use? The two different interpretations of the name would interact.

These days mobility is an important consideration, and a mobile end-system configured with the name of a resource in the private name space would anticipate a “no such domain” response when the system was relocated into the public space where the name was not delegated. Delegation of the name in the public DNS may cause an unanticipated response. Equally, the public space, and the services and resources accessible via these public DNS names would not be visible within the local scope where the name is defined in a private use context. Of course this situation poses some rather challenging policy issues in the name space. Does “squatting” on a name in a private use context confer any rights on tenure of the public name? Should the public name space avoid all names used in private contexts? Given the uncoordinated use of names in private contexts is any form of common regulation of the name space even possible in this context?

Non-DNS Names

The second example is the name space that is associated with non-DNS resolution mechanisms. One of these mechanisms is *multicast DNS* (mDNS), defined in RFC 6761^[10], which replaces the conventional unicast DNS query to a specific DNS resolver with a multicast group query, directed to the link local multicast address (224.0.0.251 or ff02::fb). All members of the multicast group receive the query and the holder of the queried name can identify itself in a multicast answer. All members of the group can learn the answer in this manner. In addition to the change of the resolution mechanism from unicast to link local multicast, RFC 6762^[11], requested that the IETF (not ICANN) reserve the generic top-level domain **.local** for use by mDNS, and thereby prevent ICANN from making a conventional unicast global public DNS delegation of the same top-level name. A related specification, *Link Local Multicast Name Resolution*, defined in RFC 4795^[12] using the Multicast group address of 224.0.0.252 and ff02::1:3, elected not to define an associated name space, so the mDNS approach was unique in some respects.

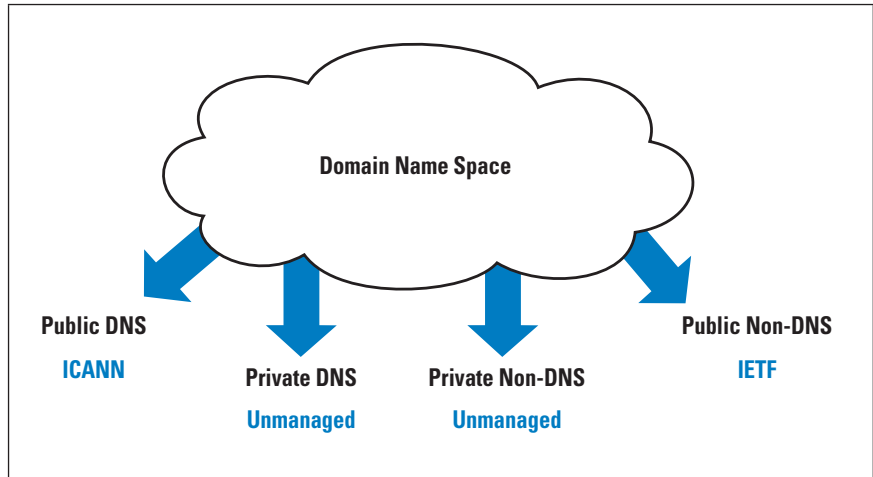
Another approach of non-DNS use of names in the domain space is the *The Onion Ring's* (TOR) use of names in the **.onion** space. Here the names within the **.onion** name space are in effect the base 32 encoded version of the public key of a defined service point, and the TOR-defined Service Directory servers are capable of performing a mapping from an encoded public key (the **.onion** name) and the desired service address. These names are not directly resolved by the DNS and connection requests for **.onion** services need to be passed into the TOR network space for resolution, RFC 7686^[13].

A third name falls into this category, and it predates the other two names by many years. The name **.localhost** refers to the local systems without further recourse to any name resolution process. It is the canonical name used to refer to oneself in the name system.

The Domain Name Space

The overall result of this process of drawing names for use out of the overall domain name space, and the entities that have some level of purview over this process is shown in Figure 1.

Figure 1: Domain Name Space Delegations



The ICANN process views the domain name space as a public good in an economic sense, and uses monetization as an intrinsic component of the name allocation function. In theory, the name space is accessible to those with an exploitation model that can recoup of expenses of acquisition of the name. In practice, the name space is accessible to those with the means to purchase a name, and there is no particular assurance that any of these names will be used in a public context. While second level names are pretty much universally accessible in `.com` or `.net`, for example, the same is probably not the case for `.google`. What was a relatively uniform common public space at the apex level of the delegated name space is now being fenced into a number of realms, many of which are private.

The publication of RFC 6761^[10] by the IETF in February 2013 essentially opened up a competing and uncoordinated channel for drawing of top-level domain names from the domain name pool. In publishing this document the IETF took what was until then a relatively static view of reserved DNS names as described in RFC 2606^[14] in 1999, and replaced it with a process that reopened up the IETF-managed name registry, using the criteria that:

“If a domain name has special properties that affect the way hardware and software implementations handle the name, that apply universally regardless of what network the implementation may be connected to, then that domain name may be a candidate for having the IETF declare it to be a Special-Use Domain Name and specify what special treatment implementations should give to that name.”
[RFC 6761]

This action is effectively unilaterally rephrasing (or “recanting”) the agreement expressed in RFC 2860 and re-defining it to mean that ICANN has purview of only those domain names that use the DNS resolution protocol, and that if the domain name uses a name resolution mechanism that does not rely on this protocol, then the name can be assigned by the IETF, via the IETF publication process. Evidently there is a set of names that are queued up to be listed using this IETF process instead of undertaking the ICANN new gTLD path^[15]. These include **.bit** (using *namecoin* resolution), **.exit** (another TOR-related name), **.gnu** and **.zkey** (using *GNU Name System* resolution), **.i2p**, **.tor** and **.carrots**.

In addition to these two parallel channels of name assignment, the private use activity continues, and names are co-opted into local use domains without any degree of effective coordination.

Clearly this story does not look good. The existence of numerous of uncoordinated activities all drawing out names from a common domain name pool is not a stable situation, nor is it in the interests of the Internet’s users. How is a user to know that names drawn from **.bit** are to be resolved using a namecoin resolution mechanism, whereas names in **.bi** or **.bid** are to be resolved using the DNS resolution protocol?

Differentiating Names?

Are there better ways to signal the resolution protocol that should be applied to a name using some additional signalling?

Should we be thinking about using a *Uniform Resource Identifier* (URI)-like syntax and using distinct schemes, such as **DNS:www.example.com** and **GNS:test.gnu**? Or using a “selector” field in a URI and using URIs of the form: **http:/namecoin/namecoin-string**?

Alternatively, we could try to push these alternate names into a single distinguished gTLD, such as **.alt**, and allow the registrars for **.alt** to register such non-DNS names in a single location in the DNS name space^[16].

We could borrow a technique used by *Internationalized Domain Names* (IDNs) and use a common prefix to denote a non-DNS name, in the same way that the character string prefix “**xn--**” denotes that the following parts of this label require pre-processing in order to produce the equivalent Unicode string. This possibility would imply that all other name forms would form part of a single name space with a single name resolution protocol, while the exception space would be clearly denoted by such a distinguished name prefix, such as, hypothetically, **.xs--gnu** for *Gnu Name System* names and **.xs--bit** names, and so on.

Behind these approaches lies a common question: What are these alternate name forms and name resolution protocols really addressing? What is the underlying issue here? If they are addressing shortfalls in the DNS, such as its lack of privacy for example, then is the appropriate answer one that includes the use of a parallel alternative name resolution protocol, or should we be looking towards the evolution of the DNS protocol to accommodate these emerging requirements? If they are addressing the ICANN position that has apparently monetized the gTLD name space and thereby blocked various other interests from accessing a gTLD name, then is the most appropriate measure for the IETF to set up of a parallel name allocation mechanism? Should the names community within ICANN undertake some deeper introspection and examine whether the gTLD program is actually catering for the full spectrum of interests in securing names for their various needs?

One Name Space?

What may be useful here is the observation that this is not a unique problem.

The radio spectrum has gone through the same process a number of times during its 100-year history, looking at the competing interests wanting access to the radio-frequency spectrum. The current spectrum allocation model contains a mix of exclusive use access arrangements. There are commercial exploitation models where actors bid for exclusive use licenses and public interest allocation models where various public sector agencies are assigned spectrum space. There are public interest and scientific use allocations, such as those used by emergency services and radio astronomers. There are also unlicensed radio spectrum allocations where there is no arrangements for exclusive access, such as are used by WiFi systems. Although a national spectrum management body is not raising revenue from these unlicensed allocations, the economic benefits of WiFi are doubtless substantial, and there is a net benefit to national economies in having this diversity of spectrum access models. The insight here is the admission that the common pool of radio spectrum space does not necessarily admit to a single exploitative model of exclusive access arrangements, and allowing a diversity of models, including that of unlicensed access, has proved to be a useful framework.

What is evident is that ICANN's gTLD process has evidently not encompassed the plurality of demand for domain names. One characterization of the outcomes of the policy for new gTLDs is that it has encouraged competitive access within a relatively narrow model of use. Access to further gTLDs within this process has many barriers, including not inconsiderable financial outlays and process overheads. The reaction has been for numerous parties to look to the IETF's management of the Special-Use Domain Name registry as an alternate means of reserving a domain name and precluding it from being used by ICANN's new gTLD program.

The rationale for entry into the IETF's Special Use Names registry—namely that the name in question uses a non-DNS resolution protocol—could be argued to be a superficial artifice that hides a more significant issue about the broad variety of the natural demand for use of names drawn from the common pool of the domain name space, and the consequent pressure for a range of means for such demands to be satisfied.

There is no doubt that the Internet's users benefit from a single coherent name space. There is considerable benefit in having the same 'name' encompass the same semantic intent and thereby 'name' the same set of services irrespective of the context, locale or time of use of that name. At the same time, the underlying technologies of name resolution, including not only the DNS resolution protocol but also other forms and means of name resolution, are subject to evolutionary pressures. It is valuable to have a means to expose these exploratory efforts in an environment of scale of use, and clearly the IETF has a role to play here. But the current mechanism of having these two bodies making uncoordinated allocations from a common name pool is not an ideal situation.

What leads to some level of unease here about the coherence of the name space is the radically different processes of evaluation of the name itself.

In the ICANN case the new gTLD process requires evaluation of the name to ensure that does not unduly infringe on existing name use, including consideration of existing brands and trademarks, designation of origin and geographic terms, issues of consumer protection through consideration of name similarity and forms of intentional passing off, potential clashes with names used by recognized international organizations, offensive terms, potential name collisions and similar environmental concerns. One could argue with the effectiveness of the process used by ICANN to evaluate these considerations, but there is some merit in the intent of ensuring that there is a process that is mindful of the larger environment of name use when considering adding a further name into this pool of use by the Internet.

The IETF's evaluation process described in RFC 6761 for admission to the Special Use Names Registry appears to admit no such similar consideration. The seven questions posed in RFC 6761 are concerned primarily with the impact of this hidden "switch" that directs applications, name resolvers, and users to understand that this name is not to be resolved by the DNS. None of these questions are concerned with the name itself, and the consequent concern is that this process could be readily abused to legitimate name squatting, and be the source of various forms of name collision. For example, the reservation of the label `.local` in the Special Use Names Registry collides with extensive conventional DNS use in local contexts^[9].

There is no record of any evaluation by the IETF of the consequences of a registration of **.local** as a reserved name for use in non-DNS contexts, with its implicit switch to a different resolution protocol of multicast DNS, that collides with pre-existing use of **.local** names in conventional local DNS contexts. Nor was there any evidence that there was consideration given to mobile users who may move in and out of environments where names in **.local** have entirely different properties and meanings, and the security issues that could result from such confusion.

However, this situation is definitely not a case of “ICANN good, IETF bad!” Far from it! But it does illustrate that there is much more to a name than might appear at the outset. The name space is indeed larger than just the DNS name resolution protocol, and this is perhaps something for ICANN to consider. At the same time names exist in a larger context of social and technical use, and this is something for the IETF to consider if it wishes to accept further reservations in the Special Use Name registry. There is also the consideration of the larger issue of whether implicit (and largely invisible) name-triggered resolution protocol switches are really in the best interests of Internet users. And for those vendors and network administrators looking for local use names to support various form of plug and play, there is the consideration of name collisions and the potential security concerns for unsuspecting users when end systems move in and out of local environments where certain name forms take on altered meanings and altered contexts of use.

If we think that a coherent and consistent name space for the Internet still has some intrinsic value, then we simply have to make some changes here to allow for a broader diversity of name use for the Internet. At the same time we must avoid stomping wilfully on each other’s toes!

References and Annotated Reading List

This is a topic that has been considered for some time, and at some considerable length, so there is no shortage of material on the topic of the name space of the Internet.

[1] Special Use Domain Names:

<http://www.iana.org/assignments/special-use-domain-names/special-use-domain-names.xhtml>

[2] Jon Postel and Joyce Reynolds, “Domain Requirements,” RFC 920, October 1984. *An early description of the DNS and its intended use.*

[3] Paul Mockapetris, “Domain Names – Concepts and Facilities,” RFC 1034, November 1987. *The original canonical specification of the DNS.*

- [4] Jon Postel, “Domain Name System Structure and Delegation,” RFC 1591, March 1994. *A description of the structure of the Domain Name space as of the early 1990’s and the description of roles and responsibilities of domain name managers.*
- [5] IAHC: <https://en.wikipedia.org/wiki/IAHC>
- [6] AlterNIC: <https://en.wikipedia.org/wiki/AlterNIC>
The Wikipedia account of the AlterNIC episode of the 1990s.
- [7] Internet Architecture Board, “IAB Technical Comment on the Unique DNS Root,” RFC 2826, May 2000. *The Internet Architecture Board’s comment on one name space and its utility for the Internet.*
- [8] Brian Carpenter, Fred Baker, and Mike Roberts, “Memorandum of Understanding Concerning the Technical Work of the Internet Assigned Numbers Authority,” RFC 2860, June 2000. *The agreement between the IETF and ICANN over names and addresses.*
- [9] Geoff Huston and George Michaelson, “On Queries to the Root,” <http://www.potaroo.net/presentations/2014-06-24-namecollide.pdf>
- [10] Stuart Cheshire and Marc Krochmal, “Multicast DNS,” RFC 6762 February 2013. *The specification of the Multicast DNS resolution protocol and the reservation of .local in the Special Use Names Registry.*
- [11] Stuart Cheshire and Marc Krochmal, “Special Use Domain Names,” RFC 6761, February 2013. *The document that re-opened the special use name registry using a rationale of non-DNS resolution technology.*
- [12] Levon Esibov, Dave Thaler, and Bernard Aboba, “Link-Local Multicast Name Resolution (LLMNR),” RFC 4795, January 2007. *A non-DNS locally scoped name resolution protocol specification.*
- [13] Jacob Appelbaum, Alec Muffet, “The “.onion” Special-Use Domain Name,” RFC 7686, October 2015. *The justification for listing .onion in the Special Use Names register.*
- [14] Donald E. Eastlake 3rd and Aliza R. Panitz, “Reserved Top Level DNS Names,” RFC 2606, June 1999.
- [15] Christian Grothoff, “Special Use Domain Names of P2P Systems,” <https://www.ietf.org/proceedings/93/slides/slides-93-dnsop-5.pdf>

- [16] Warren Kumari and Andrew Sullivan, “The ALT Special Use Top Level Domain,” Internet Draft, work in progress, September 2015, <https://tools.ietf.org/html/draft-ietf-dnsop-alt-tld-03>, *A working draft in the DNS Operations Working Group of the IETF that considers the use of .alt as a common top-level domain for Special Use contexts.*
- [17] Randall Atkinson Sally Floyd, “IAB Concerns and Recommendations Regarding Internet Research and Evolution,” RFC 3869, August 2004. *Section 3.2 considers research topics in the general area of names research.*
- [18] DNSOP Working Group – Proceedings of WG Meeting, IETF 93, July 2015, <https://www.ietf.org/proceedings/93/slides/slides-93-dnsop-5.pdf>, *A list of other names that are candidates for listing in the Special Use Names Registry.*
- [19] NSRG, “What’s In A Name: Thoughts from the NSRG,” abandoned work, September 2003, <https://tools.ietf.org/html/draft-irtf-nsrg-report-10>, *A report from the Name-space Research Group (NSRG) of the Internet Research Task Force (IRTF). It appears that the report failed to achieve consensus within the group, and was never published as an RFC.*
- [20] Joe Abley, Peter Koch, and Alain Durand, “Problem Statement for the Reservation of Top-Level Domains in the Special-Use Domain Names Registry,” <https://tools.ietf.org/html/draft-adpkja-dnsop-special-names-problem-00>, *One of those rare cases where the document’s title says it all.*
- [21] ICANN, “New gTLD program,” <http://newgtlds.icann.org/en/>, *A description of ICANN’s gTLD expansion program.*
- [22] ISO 3166 Country Codes: http://www.iso.org/iso/country_codes, *This specification is the product of TC 46, a committee of the International Organization for Standardization. This body was formed in 1947 as the outcome of collaboration between ISA, the International Federation of National Standardizing Associations, and the UNSCC, the United Nations Standards Coordinating Committee. Strictly speaking ISO does not define countries nor their standard names—that is the product of the United Nations Statistics Divisions. However ISO manages the assignment of 2 and 3 letter codes to countries, recommended as a “general purpose code.”*

GEOFF HUSTON, B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for building the Internet within the Australian academic and research sector in the early 1990s. He is author of numerous Internet-related books, and was a member of the Internet Architecture Board from 1999 until 2005. He served on the Board of Trustees of the Internet Society from 1992 until 2001. At various times Geoff has worked as an Internet researcher, an ISP systems architect, and a network operator. E-mail: gih@apnic

Quality of Service and Quality of Experience: Network Design Implications

by William Stallings, Independent Consultant
Florence Agboma, BSkyB

The Internet and enterprise IP-based networks continue to see rapid growth in the volume and variety of data traffic. Cloud computing, big data, the pervasive use of mobile devices on enterprise networks, and the increasing use of video and image traffic all contribute to the increasing difficulty in maintaining satisfactory network performance. Two key tools in measuring the network performance that an enterprise desires to achieve are *Quality of Service* (QoS) and *Quality of Experience* (QoE). QoS and QoE enable the network manager to determine if the network is meeting user needs and to diagnose problem areas that require adjustment to network management and network traffic control. This article provides an overview of QoS and QoE concepts, the relationship between the two, and the design implications of a QoE/QoS architecture.

Background and Overview

Historically, the Internet and other IP-based networks provided a *best-effort* delivery service. The term “best effort” refers to the connectionless, datagram nature of the interconnected set of networks. With best effort, a packet may be lost, duplicated, delayed, or delivered out of order, and the network does not inform sender or receiver. Traditionally, a best-effort network attempts to allocate its resources with equal availability and priority to all traffic flows, with no regard for application priorities, traffic patterns and load, or customer requirements. To protect the network from congestion collapse and to guarantee that some flows do not crowd out other flows, congestion-control mechanisms were introduced; they tended to throttle traffic that consumed excessive resources. As the intensity and variety of traffic increased, various QoS mechanisms were developed, including *Integrated Services Architecture* (ISA) and *Differentiated Services* (DiffServ) (for example, see [1]). *Service-Level Agreements* (SLAs) were also developed so that the service provided to various customers was tunable and somewhat predictable. These mechanisms and services serve two purposes: (1) allocate network resources efficiently so as to maximize effective capacity; and (2) enable networks to offer customers different levels of QoS on the basis of their requirements.


QoS is an important but increasingly insufficient tool for providing network services for many of today’s high-volume applications. To meet the needs of such applications, QoS has recently been augmented with the concept of QoE, which is a subjective measure of performance as reported by the user. Unlike QoS, which can be precisely measured, QoE relies on human opinion. QoE is important particularly when dealing with multimedia applications and multimedia content delivery.

Because QoE extends the concept of QoS to more fully tailor network services and performance to customer and user needs, it is garnering increasing attention by network protocol and system designers. The management of QoE has become a crucial concept in the deployment of future successful applications, services, and products. The greatest challenges in providing QoE are developing effective methods for converting QoE features to quantitative measures and translating QoE measures to QoS measures. Whereas now it is easy to measure, monitor, and control QoS at both the networking and application layers, and at both the end system and network sides, its management is still quite intricate.

QoS Architectural Framework

The *Telecommunication Standardization Sector of the International Telecommunication Union* (ITU-T) Recommendation Y.1291 provides an overall architectural framework that relates the various elements that go into QoS provision^[2]. Figure 1 shows the relationship between these elements, which are organized into three planes: data, control, and management. This architectural framework is an excellent overview of QoS functions and their relationships and provides a useful basis for summarizing QoS.

Figure 1: Architectural Framework for QoS Support



The *Data Plane* includes those mechanisms that operate directly on flows of data. The following discussion briefly describes each mechanism in turn.

Queue Management algorithms manage the length of packet queues by dropping packets when necessary or appropriate. Active management of queues is concerned primarily with congestion avoidance. In the early days of the Internet, the queue management discipline was to drop any incoming packets when the queue was full; it was referred to as the *tail-drop* technique. This technique has many drawbacks, including^[3]:

- There is no reaction to congestion until it is necessary to drop packets, whereas a more aggressive congestion-avoidance technique would likely improve overall network performance.
- Queues tend to be close to full, causing an increase in packet delay through a network and possibly resulting in a large batch of dropped packets for bursty traffic, necessitating many packet retransmissions.
- Tail drop may allow a single connection or a few flows to monopolize queue space, preventing other connections from getting room in the queue.

One noteworthy example of queue management is *Random Early Detection* (RED)^[4]. RED drops incoming packets probabilistically based on an estimated average queue size. The probability for dropping increases as the estimated average queue size grows.

Queuing and Scheduling Algorithms, also referred to as “queuing discipline algorithms,” determine which packet to send next; they are used primarily to manage the allocation of transmission capacity among flows.

Congestion Avoidance deals with means for keeping the load of the network sufficiently under its capacity such that the network can operate at an acceptable performance level. The specific objectives are to avoid significant queuing delays and, especially, to avoid congestion collapse.

Packet Marking encompasses two distinct functions. First, packets may be marked by network edge nodes to indicate some form of QoS that the packet should receive. An example is the *DiffServ* (DS) field in the IPv4 and IPv6 packets and the *Traffic Class* field in *Multiprotocol Label Switching* (MPLS) labels^[5]. An edge node can set the values in these fields to indicate a desired QoS. Such markings may be used by intermediate nodes to provide differential treatment to incoming packets. Second, packet marking can be used to mark packets as nonconforming; such packets may be dropped later if congestion occurs.

Traffic Classification refers to the assignment of packets to a traffic class at the edge of the network. Typically, the classification entity looks at multiple fields of a packet, such as source and destination address, application payload, and QoS markings, and determines the aggregate to which the packet belongs. This classification provides network elements a method to weigh the relative importance of one packet over another in a different class. All traffic assigned to a particular flow or other aggregate can be treated similarly. The flow label in the IPv6 header can be used for traffic classification.

Traffic Policing determines whether the traffic being presented, is on a hop-by-hop basis compliant, with prenegotiated policies or contracts. Nonconforming packets may be dropped, delayed, or labeled as nonconforming. As an example, ITU-T Recommendation Y.1221^[6] recommends the use of token bucket to characterize traffic for purposes of traffic policing.

Traffic Shaping controls the rate and volume of traffic entering and transiting the network on a per-flow basis. The entity responsible for traffic shaping buffers nonconforming packets until it brings the respective aggregate in compliance with the traffic. The resulted traffic thus is not as bursty as the original and is more predictable. For example, Y.1221 recommends the use of leaky bucket and/or token bucket for traffic shaping.

The *Control Plane* is concerned with creating and managing the pathways through which user data flows. *Admission Control* determines what user traffic may enter the network. This decision may be in part determined by the QoS requirements of a data flow compared to the current resource commitment within the network. But beyond balancing QoS requests with available capacity to determine whether to accept a request, there are other considerations for admission control. Specifically, network managers and service providers must be able to monitor, control, and enforce use of network resources and services based on policies derived from criteria such as the identity of users and applications, traffic/bandwidth requirements, security considerations, and time of day or week. RFC 2753^[7] discusses such policy-related issues.

QoS Routing determines a network path that is likely to accommodate the requested QoS of a flow. This function contrasts with the philosophy of traditional routing protocols, which generally are looking for a least-cost path through the network. RFC 2386^[8] provides an overview of the issues involved in QoS routing, which is an area of ongoing study.

Resource Reservation reserves network resources on demand for delivering desired network performance to a requesting flow. The resource-reservation mechanism that has been implemented for the Internet is the *Resource Reservation Protocol* (RSVP)^[9].

The *Management Plane* contains mechanisms that affect both control- and data-plane mechanisms. It deals with the operation, administration, and management aspects of the network. A *Service-Level Agreement* (SLA) is the agreement between a customer and a provider of a service that specifies the level of availability, serviceability, performance, operation, or other attributes of the service. SLAs are discussed subsequently.

Traffic Metering and Recording concerns monitoring the dynamic properties of a traffic stream using performance metrics such as data rate and packet-loss rate. It involves observing traffic characteristics at a given network point and collecting and storing the traffic information for analysis and further action. Depending on the conformance level, a meter can invoke necessary treatment (for example, dropping or shaping) for the packet stream. A subsequent section discusses the types of metrics that are used in this function.

Traffic Restoration refers to the network response to failures. It encompasses numerous protocol layers and techniques.

Policy refers to a set of rules for administering, managing, and controlling access to network resources. The rules can be specific to the needs of the service provider or reflect the agreement between the customer and service provider, which may include reliability and availability requirements over a period of time and other QoS requirements.

Service-Level Agreements


An SLA is a contract between a network provider and a customer that defines specific aspects of the service that is to be provided. The definition is formal and typically defines quantitative thresholds that must be met. An SLA typically includes the following information:

- *A description of the nature of service to be provided:* A basic service would be IP-based network connectivity of enterprise locations plus access to the Internet. The service may include additional functions such as web hosting, maintenance of domain name servers, and operation and maintenance tasks.
- *The expected performance level of the service:* The SLA defines numerous metrics, such as delay, reliability, and availability, with numerical thresholds.
- *The process for monitoring and reporting the service level:* This function describes how performance levels are measured and reported.

The types of service parameters included in an SLA for an IP network are similar to those provided for *Frame Relay* and *Asynchronous Transfer Mode* (ATM) networks. A key difference is that, because of the unreliable datagram nature of an IP network, it is more difficult to realize tightly defined constraints on performance, compared to the connection-oriented Frame Relay and ATM networks.

Figure 2 shows a typical configuration that lends itself to an SLA. In this case, a network service provider maintains an IP-based network. A customer has many private networks (for example, LANs) at various sites. Customer networks are connected to the provider via access routers at the access points. The SLA dictates service and performance levels for traffic between access routers across the provider network. In addition, the provider network links to the Internet and thus provides Internet access for the enterprise.

Figure 2: Typical Framework for Service-Level Agreement



For example, the standard SLA provided by Cogent Communications for its backbone networks includes the following items:

- *Availability*: 100% availability
- *Latency (delay)*: Monthly average Network Latency for packets carried over the COGENT Network between Backbone Hubs for the following regions: Intra-North America: ≤ 45 ms; Intra-Europe: ≤ 35 ms; New York to London: ≤ 85 ms; Los Angeles to Tokyo: ≤ 120 ms.

Latency is defined as the average time taken for an IP packet to make a round trip between Backbone Hubs within a region. COGENT monitors aggregate latency within the COGENT Network by monitoring round-trip times between a sample of Backbone Hubs on an ongoing basis.

- *Network Packet Delivery (reliability)*: Average monthly Packet Loss no greater than 0.1% (or successful delivery of 99.9% of packets). Packet Loss is defined as the percentage of packets that are dropped between Backbone Hubs on the COGENT Network.

An SLA can be defined for the overall network service. In addition, SLAs can be defined for specific end-to-end services available across the carrier's network, such as a virtual private network, or differentiated services.

IP Performance Metrics

The *IP Performance Metrics Working Group* (IPPM) is chartered by the *Internet Engineering Task Force* (IETF) to develop standard metrics that relate to the quality, performance, and reliability of Internet data delivery. Two trends dictate the need for such a standardized measurement scheme:

- The Internet has grown and continues to grow at a dramatic rate. Its topology is increasingly complex. As its capacity has grown, the load on the Internet has grown at an even faster rate. IP-based enterprise networks have exhibited similar growth in complexity, capacity, and load. The sheer scale of these networks makes it difficult to determine quality, performance, and reliability characteristics.
- The Internet serves a large and growing number of commercial and personal users across an expanding spectrum of applications. Similarly, private networks are growing in terms of user base and range of applications. Some of these applications are sensitive to particular QoS parameters, leading users to require accurate and understandable performance metrics.

A standardized and effective set of metrics enables users and service providers to have an accurate common understanding of the performance of the Internet and private internets. Measurement data is useful for a variety of purposes, including:

- Supporting capacity planning and troubleshooting of large complex internets
- Encouraging competition by providing uniform comparison metrics across service providers
- Supporting Internet research in such areas as protocol design, congestion control, and quality of service
- Verification of service-level agreements

The metrics are defined in three stages:

- *Singleton metric*: The most elementary, or atomic, quantity that can be measured for a given performance metric. For example, for a delay metric, a singleton metric is the delay experienced by a single packet.

- *Sample metric*: A collection of singleton measurements taken during a given time period. For example, for a delay metric, a sample metric is the set of delay values for all of the measurements taken during a one-hour period.
- *Statistical metric*: A value derived from a given sample metric by computing some statistic of the values defined by the singleton metric on the sample. For example, the mean of all the one-way delay values on a sample might be defined as a statistical metric.

The measurement technique can be either *active* or *passive*.

Active techniques require injecting packets into the network for the sole purpose of measurement. This approach has several drawbacks. The load on the network is increased, in turn possibly affecting the desired result. For example, on a heavily loaded network, the injection of measurement packets can increase network delay, so that the measured delay is greater than it would be without the measurement traffic. In addition, an active measurement policy can be abused for denial-of-service attacks disguised as legitimate measurement activity.


Passive techniques observe and extract metrics from existing traffic. This approach can expose the contents of Internet traffic to unintended recipients, creating security and privacy concerns. So far, the metrics defined by the IPPM working group are all active.

For the sample metrics, the simplest technique is to take measurements at fixed time intervals, known as periodic sampling. There are several problems with this approach. First, if the traffic on the network exhibits periodic behavior, with a period that is an integer multiple of the sampling period (or vice versa), correlation effects may result in inaccurate values.

Also, the act of measurement can perturb what is being measured (for example, injecting measurement traffic into a network alters the congestion level of the network), and repeated periodic perturbations can drive a network into a state of synchronization (for example, [10]), greatly magnifying what might individually be minor effects. Accordingly, RFC 2330^[11] recommends Poisson sampling. This method uses a Poisson distribution to generate random time intervals with the desired mean value.

Figure 3 illustrates the packet-delay variation metric. This metric is used to measure jitter, or variability, in the delay of packets traversing the network. The singleton metric is defined by selecting two packet measurements and measuring the difference in the two delays. The statistical measures use the absolute values of the delays.

Figure 3: Model for Defining Packet-Delay Variation



QoS for Streaming Video

It is worthwhile to comment on the relationship between QoS concerns and streaming video, which perhaps accounts for the greatest volume of traffic on the Internet. First, consider the network transmission requirements for real-time video, such as video teleconferencing. For such applications QoS metrics such as latency and jitter are important, and they impose significant requirements on the networks through which the real-time traffic passes.

For streaming video, in contrast, such QoS requirements can be relaxed by equipping the application to handle a wider range of responses. A motivation for this approach is that QoS is not ubiquitously deployed in networks and thus is not available for all possible streams between content deliverers and content consumers. In the case of streaming video in particular, applications have come to assume that the receiving device is equipped with generous levels of memory, and the application uses this memory as a playback buffer. This system allows the application to use TCP for reliable delivery. The TCP connection can be exposed to far higher levels of network jitter and even packet loss than would be tolerable by a real-time packet stream. Thus, to a significant extent, this approach to video streaming avoids the need to use QoS as a strict precondition for streaming video over the Internet.

From QoS to QoE

The literature contains numerous different, though similar, definitions of QoE. To provide a common working definition, the EU-sponsored *European Network on Quality of Experience in Multimedia Systems and Services* (see [12]) has developed a definition of QoE that reflects broad industry and academic consensus:

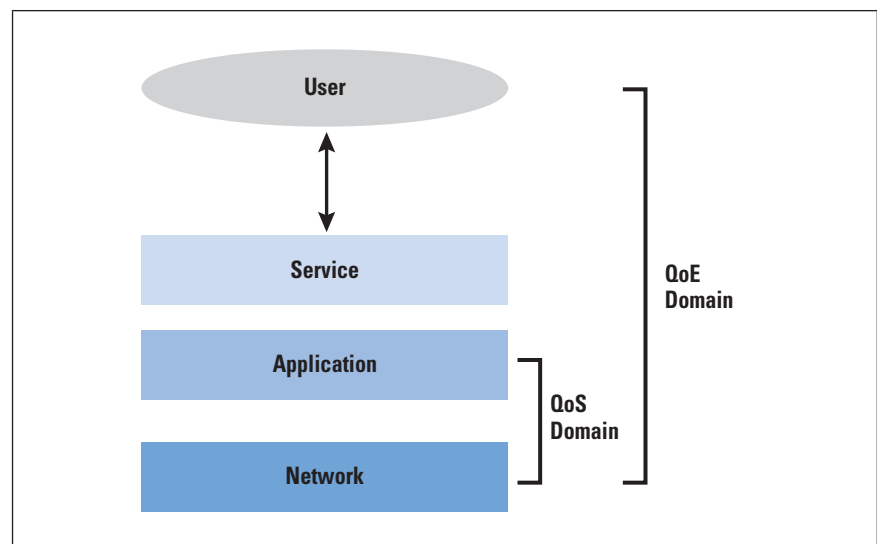
“Quality of Experience (QoE) is the degree of delight or annoyance of the user of an application or service. It results from the fulfillment of his or her expectations with respect to the utility and/or enjoyment of the application or service in the light of the user’s personality and current state.”

The QoE methods to measure this degree of delight or annoyance are discussed later in this article.

The Layers of QoE

For any type of service (for example, IPTV), multiple QoS parameters contribute to the overall user’s perception of quality. The concept of QoE in addition to QoS mechanisms has been proposed in [13], [14], [15], and [16] as a QoE/QoS layered approach by which the requirements of the users drive network-dimensioning strategies. The QoE/QoS layered approach does not substitute the QoS aspect of the network, but instead, user and service-level perspectives are both complementary as shown in Figure 4.

Figure 4: QoE/QoS Layered Model



The levels from the layered approach follow:

- *User*: The user interacts with the service. It is the user’s degree of delight or annoyance from using the service that is to be measured.
- *Service*: This level is the virtual level where the user’s experience of the overall performance of the service can be measured. It is the interface where the user interacts with the service (for example, the visual display to the user).
- *Application-level QoS (AQoS)*: This level deals with the control of application-specific parameters such as content resolution, bitrate, frame rate, colour depth, codec type, layering strategy, sampling rate, and number of channels.

- *Network-level QoS (NQoS)*: This level is concerned with the low-level network parameters such as service coverage, bandwidth, delay, throughput, and packet loss.

Background on the Layers of QoE

Being linked to human perception, QoE is hard to describe quantitatively, and it varies from person to person. The complexities of QoE at the user level stem from the differences between individual user characteristics, of which some might be of a variant nature (that is, changing often with time), whilst others are of a relatively stable nature. Examples could include gender, age, attitudes, prior experience, expectations, socio-economic status, cultural background, educational level, etc. Thus, it becomes a challenge to derive unified QoE metrics for all users and their contexts. Reiter et al.^[17] provide more details on the human factors that may influence QoE. The current practice in any QoE measurement is to identify and control the relatively stable characteristics of a user in a way that is satisfactory to at least a large proportion of the potential user group.

Multiple layers might impact the user QoE for any application. The service level provides a virtual level where the user's tolerance thresholds for a particular service could be measured. As an illustration, the QoE measures from the user perspective for streaming applications could be a) start-up time, b) audiovisual quality, c) channel-change delay, and d) buffering interruptions, whilst the QoE measures for web-browsing applications could be waiting times.

Quite often, the network capacity will dictate the bandwidth that may be used for transmission. At the application-level QoS, there might be a trade-off between quality and size. For audio, a higher sampling rate, for example, 96 kHz, might allow for more information to be perceived compared with 48 kHz, but at the expense of a bigger file size. Traditional telephony speech might be limited to 8 kHz because of the bandwidth capacity. For video, a high resolution might require more bandwidth than low resolution. There are huge varieties of device screens in all kinds of sizes, featuring varied aspect ratios. The one commonality in this array of equipment is that they are all capable of rescaling the video, as is the case for a native player going to the full-screen mode. For a given bitrate, there might be a trade-off between lower resolutions in pixels (images being slightly blurred) with fewer artifacts versus higher resolutions that provide a sharper image but possibly with more artifacts. Most compression standards might use a block-based and motion-compensation coding scheme, and as a result additional compression artifacts are added to the decoded video.

Network-level QoS parameters could impact QoE in a variety of ways.

The network delay could impact QoE, especially for interactive services. For instance, the interactive nature of web browsing that requires multiple retrieval events within a certain window of time might be affected by delay variations of the network. *Voice over IP* (VoIP) services might have the stringent response-time demands, whereas email services might tolerate much longer delays. The different distribution methods of streaming video over the network might affect QoE in different ways. HTTP-based adaptive streaming, which uses TCP, might react to bandwidth constraints and CPU capacity in two ways: either the streaming switches between streaming the different bitrate encodings depending on available resources, or a frame freeze (rebuffering) due to incoming packet starvation in the player buffer. The continuous bitrate switches and rebuffering affect QoE badly. The other distribution method, *User Datagram Protocol* (UDP), might use multicast to replicate the streams throughout the network. Quite often, a resilient coding scheme and a flow-control mechanism might be implemented to maintain the viewing experience despite the effects of poor network conditions.

For reasons that should now be apparent, the background on the layers of QoE suggests that the effect of QoE could be an attribute of only the application layer or a combination of both the application and network layers. Although the trade-offs between quality and network capacity may begin with application-level QoS due to network capacity considerations, an understanding of the user requirements at the service level (that is, in terms of QoE measures) would enable a better choice of application-level QoS parameters to be mapped onto the network-level QoS parameters. A scenario that aims at controlling QoE using QoS parameters as actuators is discussed later in this article. Taking the QoE/QoS approach as a whole entity rather than single entities might aid in providing better QoEs, and potentially could lead to a better-managed delivery infrastructure.

Key Factors That Determine QoE

The nature of QoE, which comprises many layers of interaction between the enabling elements of service delivery and the human user, makes measuring and improving user experiences a challenging task. To understand QoE we must account for both technical and nontechnical factors. Many factors contribute to producing a good QoE. Moller et al.^[18] provide useful perspectives on factors that influence QoE. Here, we discuss the following key factors that might influence QoE:

User Demographics: The context of demographics herein refers to the relatively stable characteristics of a user that might indirectly influence perception, and intimately affects other technical factors to determine QoE. The findings from [19] suggest that user demographics may influence QoE. In the context of adopting HD voice telephony, the different user groups had significantly different quality ratings.

The grouping of users was based on demographic characteristics such as their attitudes towards adoption of new technologies, socio-demographic information, socio-economic status, and prior knowledge. Cultural background is another user demographic factor that might also influence perception due to cultural attitude to quality^[20].

Type of Device: Different device types possess varying characteristics that may affect QoE. An application designed to run on more than one device type, for example on a connected TV device such as Roku and on an iOS device such as an iPhone, may not deliver the same QoE on every device.

Content: The content being distributed via Internet delivery can range from interactive content specifically curated according to personal interests to content that is produced for linear TV transmission. The different characteristics of the video might require different system properties in terms of the quality being produced. According to [21], people tend to watch *Video on-Demand* (VoD) content with a higher level of engagement than its competing alternative, linear TV. This trend may be because users make an active decision to watch a particular VoD content, and as a result, give their full attention to it. One could infer that for VoD, users might be less tolerant of any quality degradations because of their high level of engagement.

Connection Type: The type of connection used to access the service influences users' expectations and their QoEs. Users have been found to have lower expectations when using 3G in contrast to a wireline connection when in fact both connection types were identical in terms of their technical conditions^[22]. Agboma et al.^[23] found users' expectations to be considerably lowered and more tolerant to visual impairments on small devices. Conventional QoS management practices cannot account for these psychological factors.

Media (audiovisual) Quality: This factor may be observed as a significant one affecting QoE, because it is the part of a service that users notice most. The integration of audio and video quality appears to be content-dependent. For less-complex scenes (for example, head and shoulder content), audio quality is slightly more important than video quality. In contrast, for high-motion content, video quality has been found to be significantly more important than audio quality^[24]. Other studies^[25] suggest that the optimum audio/video bitrate allocation depends on scene complexity. For instance, visually complex scenes would benefit from the allocation of higher bitrates, with relatively more bits allocated towards audio, because high-audio bitrates seem to produce the best overall audiovisual quality.

Network: Content delivery via the open Internet is highly susceptible to the effects of delays, jitter, packet loss, and available bandwidth. For users, delay and jitter cause frame freeze and the lack of lip synchronization between what is heard (audio) and what is seen (video). Content delivery is guaranteed using a TCP/IP delivery mechanism. However, under bad network conditions, the frequency of rebuffering, and the implementation of a video player heuristics, might affect QoE. Rebuffering interruptions in IP video playback is considered the worst degradation on user QoE and should be avoided at the cost of startup delay^[26]. On the same note, QoE for a given startup delay strongly depends on the concrete application context and user expectations^[27]. In spite of the different QoE factors that are concerned with the network, reliability and a strong wireless signal is crucial for consuming TV-like services anytime, anywhere, and from any device.

Usability: Another QoE factor is the amount of effort that is required to navigate through the service. The design should render good quality without a great deal of technical input from the user before or after service consumption.

Cost: The long-established practice of judging quality by price implies that expectations are price-dependent. If the tariff for a certain service quality is high, users may be highly sensitive to any quality degradations. A scenario of QoE-based charging is demonstrated in [28] and [29] to analyze a situation where price is used to reflect the value of a service, and at the same time, part of the user-context factor. While [28] offers a trade-off between user expectations and price, its deployment may yield unexpected complexities.

QoE Measurement Methods

QoE measurement techniques evolved through the adaptation and application of psychophysics methods during the early stages of television systems (See [30] for more details). Here, we introduce three QoE measurement methods: subjective assessment methods, objective assessment methods, and end-user device analytics as an alternative to measure QoE. Hereafter, streaming video will be the focus of discussion.

Subjective Assessment

In *Subjective Assessment* of QoE, experiments are carefully designed to a high level of control (such as in a controlled laboratory, field tests, or crowdsourcing environments) so that the validity and reliability of the results can be trusted. It might be useful to consult expert advice during the initial design of the subjective experiment, because the topics of experimental design, experimental execution, and statistical analysis are complex. The different phases discussed in the following paragraphs are an abstract. A methodology to obtain subjective QoE data might consist of the following phases:

Characterize the Service: The task at this stage is to choose the QoE measures that affect the users' experience the most. As an example, consider a multimedia conferencing service; previous studies have shown that the quality of the voice takes precedence over the quality of video^[31]. Also, the video quality required for such applications does not demand a very high frame rate, provided that audio-to-video synchronization is maintained. Thus, the resolution of individual frames can be considerably lower than the case of other video streaming services, especially when the size of the screen is small (such as that of a mobile phone). Therefore, in multimedia conferencing, the QoE measures might be prioritized as voice quality, audio-video synchronization, and image quality.

Design and Define Test Matrix: Once the service has been characterized, the QoS factors that affect the QoE measures can be identified. For instance, the video quality in streaming services might be directly affected by network parameters such as bandwidth, packet loss, and encoding parameters such as frame rate, resolution, and codec. The capability of the rendering device will also play a significant role in terms of screen size and processing power. However, testing such a large combination of parameters may not be feasible. This draft matrix could be reduced to more achievable realistic test conditions by eliminating the combinations that have similar effects on QoE.

Specify Test Equipment and Materials: Subjective tests should be designed to specify test equipment that will allow controlled enforcement of the test matrix. For instance, to assess the correlation between NQoS parameters and the perceived QoE in a streaming application, at least a client device and a streaming server separated by an emulated network are needed. If the objective is to evaluate how different device capabilities impact QoE, then a video content is chosen to produce formats that can run in each of the client devices under scrutiny.

Identify Sample Population: A representative sample population is identified, possibly covering different classes of users categorized by the user demographics that are of interest to the experimenter. Depending on the target environment for the subjective test, at least 24 test subjects (that is, participants) have been suggested for a controlled environment (for example, a laboratory) and at least 35 test subjects for a public environment. Fewer subjects may be used for pilot studies to indicate trending^[32]. The use of *crowdsourcing* in the context of subjective assessment is still nascent, but it has the potential to further increase the size of the sample population and could reduce the completion time of the subjective test.

Subjective Methods: The recommendations include several subjective assessment methodologies, for example [33], [34], and [35]. A recent recommendation^[32] addresses the context of Internet video and distribution of quality television in any environment. These recommendations provide guidelines on topics such as the different subjective test designs, rating scales, and experiment durations.

There has also been some interest in developing alternative subjective methodologies for time-varying system characteristics, see [36] and [37]. Typically, each test subject is presented with the test conditions under scrutiny along with a set of rating scales that allows the correlation of the users' responses with the actual QoS test conditions being tested. There are several rating scales, depending on the design of the experiment. Other scale methods can be found in [32], and acceptance methods in [38].

Analysis of Results: When the test subjects have rated all QoS test conditions, a post-screening process might be applied to the data to remove any erroneous data from a test subject who appears to have voted randomly. Depending on the design of the experiment, a variety of statistical approaches could be used to analyze results. For the sake of brevity, statistical analysis of the results is outside the scope of this article. The most common quantification method is the *Mean Opinion Score* (MOS), which is the average of the opinions collected for a particular QoS test condition. The results from subjective assessment experiments are used to quantify QoE, and to model the impacts of QoS factors. Other authors^[39] have gone beyond deriving QoS MOS functions to QoE management. The rationale and limitations of MOS are discussed later in this article.

Subjective experiments require significant planning and design so as to produce reliable subjective MOS ratings. However, they are time-consuming and expensive to carry out and are not feasible for real-time in-service monitoring. Therefore, the use of objective models is often desirable.

Objective Assessment

In *Objective Assessment* of QoE, computational algorithms provide estimates of audio, video, and audiovisual quality as perceived by the user. Each objective model targets a specific service type. For example, ITU-T P.1201 predicts the impact of IP network impairments for IPTV applications and multimedia streaming applications^[40]. ITU-T J.341 targets video-quality measurement in *High-Definition Television* (HDTV) noninteractive applications^[41]. Other proprietary objective models do exist. For a given scope of quality features, the goal of any objective model is to find the optimum fit that strongly correlates with data obtained from subjective experiments. The following phases presented here should not be considered exhaustive, but they are meant to illustrate a process of obtaining objective QoE data. A methodology to obtain objective QoE data might consist of the following phases:

Database of Subjective Data: A starting point might be the collection of a group of subjective datasets as this list could serve as benchmark for training and verifying the performance of the objective model. A typical example of one of these datasets might be the subjective QoE data generated from well-established subjective testing procedures, as discussed earlier. The selection of the subjective datasets should typically reflect the use cases of the objective model.

Preparation of Objective Data: The data preparation for the objective model might typically include a combination of the same QoS test conditions as found in the subjective datasets, as well as other complex QoS conditions. A variety of pre-processing procedures might be applied to the video data prior to training and refinement of the algorithm.

Objective Methods: Various algorithms in the literature could provide estimates of audio, video, and audiovisual quality as perceived by the user. Some algorithms might be developed, and test subjects trained, for a specific perceived quality artifact, whilst other algorithms and test subjects might be trained for a wider scope of perceived quality artifacts. Examples of the perceived artifacts might include blurring, blockiness, unnatural motion, pausing, skipping, rebuffering, and imperfect error concealment after transmission errors. See [42] for a good overview, and for the development of objective video-quality prediction.

Verification of Results: After the objective algorithm has processed all QoS test conditions, the predicted values might benefit from a post-screening process to remove any outliers, the same concept as applied to the subjective datasets. The predicted values from the objective algorithm might be in a different dimension compared to the subjective QoE datasets. The predicted values might be transformed to the same scale as obtained in the subjective experiments (for example, into the MOSs) to allow for a linear comparison, and so that the optimum fit between the predicted QoE values and subjective QoE data can be obtained. This transformation might be an integral module of the objective model. The statistical analysis that might be applied to calibrate the scale of an objective model is outside the scope of this article.

Validation of Objective Model: The objective data analysis might be evaluated with respect to its prediction accuracy, consistency, and linearity by using different subjective datasets. It may also be worth noting that the performance of the model might depend on the training datasets and the verification procedures. See [43] for more details on calibrating and validating objective models. The *Video Quality Experts Group* (VQEG) validates the performance of objective perceptual models that result in ITU recommendations and standards for objective quality models for both television and multimedia applications^[44]. The practical deployment of such objective models is discussed in the background of QoE measurement methods.

End-User Device Analytics

End-user device analytics could provide an alternative to measuring QoE as experienced by the end user. Real-time data such as the connection time, bytes sent, and average playback rate are collected by the video player application for each video viewing session and fed back to a server module where the data is pre-aggregated and then turned into actionable QoE measures.

Some of the metrics reported for per-user and aggregate viewing sessions include startup delay, rebuffering delays, average bitrates, and frequency of bitrates switches. Operators might be more inclined to associate viewers' engagement to QoE because better QoEs might make viewers less likely to abandon a viewing session, leading to increased monetization of assets and low subscriber turnover.

The definition of viewer engagement may have different meanings for different operators and contexts. First, operators might like to know which viewer engagement metrics affect QoE the most in order to guide the design of the delivery infrastructures. Secondly, they might also like to quickly identify and resolve service outages and other quality issues. A minute of encoder glitch could replicate throughout the *Internet Service Providers* (ISPs) and the various delivery infrastructures, and affect all their customers. Operators might like to know the scale of this impact, and how it affects users' engagement. The cost of getting the viewer experience wrong often makes the news headlines^[45]. Finally, they would like to understand their customers' demographics (connection methods, type of device, and bitrates of the consumed asset) within a demographic region so that they can fully monetize their assets, and use their other resources strategically.

QoE enthusiasts advocate QoE measurement to be a multidisciplinary approach that seeks to explain its findings, building on general laws of perception, sociology, and user psychology^[39]. If the end-user device analytics is used as a means of QoE measurement, many unexplained variables may not be accounted for (for example, why a user exits a service). For instance, the viewer's tolerance thresholds of QoS parameters for live and VoD might have different QoE patterns, and another dimension of complexity to include might be premium vs. free content. Also, a lack of interest in watching the content, not necessarily an effect of poor QoE, might make a user exit a service. Viewer engagement metrics is measured objectively bypassing subjective studies and surveys.

Some attempts at addressing these unexplained variables can be found in the literature. For example, [46] used the fraction of video viewed as a measure of engagement because this factor can be measured objectively. The data that appeared to belong to early quitters were systematically removed from their analysis to provide a clearer understanding of how the QoE measures affected viewer engagement. A slightly different approach to measuring viewer engagement can be found in [47]. However, with a huge database of aggregated data in the range of millions of viewing session logs, some profile classes of how the QoE measures affect viewer engagement might emerge.

Background on QoE Measurement Methods

The MOS appears to be the *de facto* standard metric for QoE. The possible reasons could be its long-term establishment in telephony networks, perhaps its widespread acceptance on the merits that it can be easily understood, and a metric for benchmarks. There are different types of MOS values, and different test methodologies to produce them, (see [32] for more details). Figure 5 shows the five-point absolute category rating MOS scale that is commonly used.

Figure 5: Five-Point Rating Scale

Score	Label
5	Excellent
4	Good
3	Fair
2	Poor
1	Bad

The MOS value is the average opinion for a given QoS tests condition, not necessarily for the individual users because different users have different opinions. Additional information such as statistical uncertainty in terms of confidence intervals is usually encouraged. The MOS is considered to be characteristic of only the experiment and the group of test subjects from which it was derived.

While there is a reference methodology to produce MOS, it has to be interpreted within context. First, the MOS value obtained for a particular QoS test condition in a subjective experiment may depend on the range of the QoS test conditions used in the experiment. This dependence might be due to test subjects who re-calibrate their use of the rating scale to the conditions in the experiment. An appropriately designed experiment that has a practice period at the start of the experiment, and the test conditions include the best and worst conditions, minimizes the effects of the aforementioned behavior.

Secondly, direct comparisons of MOS scores obtained from separate experiments are generally not meaningful. They are meaningful only if the experiments have been specially designed to enable such comparisons. Data from such specially configured experiments must be studied and shown that their MOS comparisons are statistically valid. Biases in the rating scale interpretation might exist such as differences in interpretation and use of rating scales across cultures; test subject profile, for example, age and technology exposure; test environment; and the presentation order of the test conditions^[48].

Thirdly, it is possible that different objective models that have been trained and optimized using different subjective contexts will predict non-identical MOS values for the same QoS conditions. Objective models are usually developed and optimized for a specific scope of quality features. As a consequence, comparisons between MOS predictions and thresholds can be reliably made only if the thresholds are chosen in the context of the MOS model. See [46] for MOS interpretation and reporting.

The ITU standardization activities^[49] classify objective quality assessment methods into five main categories: media-layer models, parametric packet-layer models, parametric planning models, bitstream-layer models, and hybrid models. Depending on the input parameters to the model and the specific service type (for example, speech, audio, video, and multimedia), each category is aimed at predicting QoE.

The practical deployment of such objective models might benefit in-service monitoring positioned at any one node, or at all nodes, within the content-delivery ecosystem. The node(s) could be at the headend incoming feeds, distribution networks, and at locations of endpoints that are customer equipment. For example, the media-layer model (also known as the perceptual model) might be best suited for quality assurance at the headend, and for benchmark purposes. The media-layer model uses both the source video and the impaired video to predict QoE. The parametric-packet layer model uses only packet-header information such as IP packets, for example, UDP, and extracts the required information to predict QoE. It does not consider encoding distortions, but owing to its light-weight implementation it might be appropriate for in-service monitoring distribution networks. On the other hand, the bitstream-layer models might be appropriate for locations of endpoints that are customer equipment, because they analyze up to the bitstream level. This model in turn does not consider the decoded (impaired) output. The perceptual hybrid model, which combines the media layer and bitstream layer, might use both the bitstream information and the decoded output to predict QoE. While objective assessment seems to offer real-time QoE measurements, some categories might not meet industry demands. Hence, end-user device analytics as a method to QoE measurement appears to be an alternative approach.

Currently, there is the lack of a reference methodology for end-user device analytics as a method of QoE measurement, analogous to MOSs found in subjective assessments^[32] and objective assessments^[42]. A limiting factor to this development might be the restricted rights governing service providers, or the likes, on the usage of such databases. This situation makes it challenging for researchers, service providers, and delivery infrastructures to focus their efforts on developing better delivery infrastructures.

Subjective experiments could still be the most accurate way to measure perceived QoE, and the only way to obtain reliable ground-truth data used in benchmarking objective QoE models.

Linking QoS to QoE

A first glance indicates a considerable mismatch between the concepts of QoS and QoE. This mismatch can be seen in the ITU-T definitions of the two terms. QoS is defined in ITU-T E.800^[50] as the “totality of characteristics of a telecommunications service that bear on its ability to satisfy stated and implied needs of the user of the service.” These characteristics are objective measures consisting of quantitative variables and attributes that may be present or absent. Thus, QoS characteristics are objective and can be objectively measured. By contrast, ITU-T P.10^[51] defines QoE as “the overall acceptability of an application or service, as perceived subjectively by the end-user.” Thus, QoE is called out specifically as consisting of subjective measures.

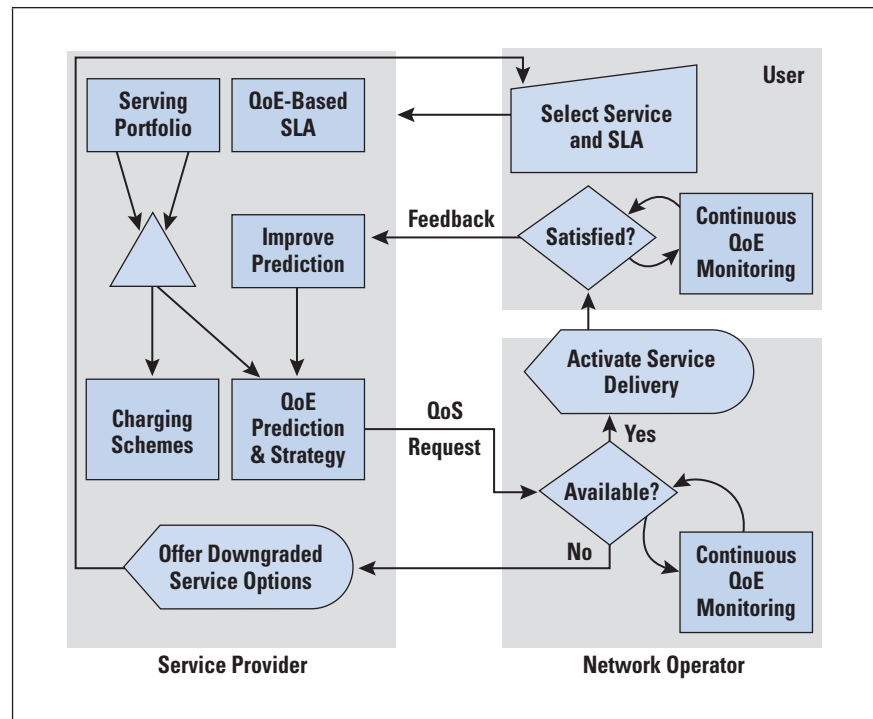
Both network service providers and customers have become accustomed to developing SLAs based on QoS measures. Recent years have seen a growing awareness on the part of both providers and customers that QoE is the more important concept, and that ways must be found to tie network performance parameters, as committed in an SLA, to the desired user QoE. For a typical network and service environment for a particular customer, numerous QoS metrics will impact overall QoE. The focus of ongoing research and product development is on developing reliable techniques, acceptable to both service provider and customer, for correlating QoS performance metrics with QoE as measured by MOS. ITU-T G.1080 identifies two ways in which such correlations can be exploited:

- Given a QoS measurement, predict the expected QoE for a user, with appropriate assumptions.
- Given a target QoE for a user, deduce the required network service performance, with appropriate assumptions.

Service providers can take the first approach and provide a range of QoS offerings with an outline of the QoE that their customers might reasonably expect. Customers can take the second approach by defining the required QoE and then determining what level of service will meet their needs.

Figure 6 illustrates a scenario for the second approach, where the user can make a selection from a range of services, including the required *level of service* (SLA). By contrast to the purely QoS-based management, the SLA here is not expressed in terms of raw network parameters. Instead, the user indicates a QoE target; it is the service provider that maps this QoE target together with the type of service selected, onto QoS demands.

Figure 6: User-Centric Service Delivery



For instance, in the case of multimedia streaming service, the user may simply choose between two QoE levels (high or low). The service provider selects the appropriate quality-prediction model and management strategy (for example, minimize network resource consumption) and forwards a QoS request to the operator. It is possible that the network cannot sustain the required level of QoS, making it impossible to deliver the requested QoE. This situation leads to a signal back to the user, prompting a reduced set of services/QoE values.

Assuming that the network can support the service, delivery can be activated. During service operation, two monitoring and control loops run concurrently: one at network level and the other at service level. The latter allows the user to switch to a different level of QoE (for example, to get a cheaper service or to request higher quality). If the user generates no explicit feedback, it means that the user is satisfied, confirming that the quality-prediction model is working. In this way, the quality-prediction model continues to be redefined during service delivery, allowing it to evolve as user needs and devices change over time.

There are so many varied components with this approach, which extends from the management complexity of a *Content-Delivery Network* (CDN) service based on QoE SLAs to service billing. However, focusing on integrating QoE as part of the management methodology during the design and development of services ensures a user-centric perspective, and helps to move beyond the current network design principles.

References

- [1] Geoff Huston, “QoS—Fact or Fiction?” *The Internet Protocol Journal*, Volume 3, No. 1, March 2000.
- [2] ITU-T, “An Architectural Framework for Support of Quality of Service in Packet Networks,” Recommendation Y.1291, May 2004.
- [3] Sally Floyd and Van Jacobson, “Random Early Detection Gateways for Congestion Avoidance,” *IEEE/ACM Transactions on Networking*, August 1993.
- [4] Jon Crowcroft, Bob Braden, Steve Deering, Sally Floyd, Bruce Davie, Van Jacobson, and Deborah Estrin, “Recommendations on Queue Management and Congestion Avoidance in the Internet,” RFC 2309, April 1998.
- [5] William Stallings, “MPLS,” *The Internet Protocol Journal*, Volume 4, No. 3, September 2001.
- [6] ITU-T, “Traffic Control and Congestion Control in IP-based Networks,” Recommendation Y.1221, June 2010.
- [7] Dimitrios Pendarakis, Raj Yavatkar, and Roch Guerin, “A Framework for Policy-based Admission Control,” RFC 2753, January 2000.
- [8] Eric S. Crawley, Raj Nair, Bala Rajagopalan, and Hal Sandick, “A Framework for QoS-based Routing in the Internet,” RFC 2386, August 1998.
- [9] Lixia Zhang, Steve Berson, Shai Herzog, and Sugih Jamin, “Resource ReSerVation Protocol (RSVP) – Version 1 Functional Specification,” RFC 2205, September 1997.
- [10] Sally Floyd, “TCP and Explicit Congestion Notification,” *ACM Computer Communication Review*, October 1994.
- [11] Vern Paxson, Guy Almes, Jamshid Mahdavi, and Matt Mathis, “Framework for IP Performance Metrics,” RFC 2330, May 1998.
- [12] Sebastian Möller, Patrick Le Callet, and Andrew Perkis, “Qualinet White Paper on Definitions on Quality of Experience – Output Version of the Dagstuhl Seminar 12181,” *European Network on Quality of Experience in Multimedia Systems and Services*, COST Action IC 1003, 2012.
- [13] Florence Agboma and Antonio Liotta, “QoE for Mobile TV Services,” In *Multimedia Transcoding in Mobile and Wireless Networks*, IGI Global, 2009.

- [14] Mario Siller, “An Agent-Based Platform to Map Quality of Service to Experience in Active and Conventional Networks,” University of Essex, 2006.
- [15] Milos Ljubojević et al., “Influence of Resolution and Frame Rate on the Linear in-Stream Video Ad QoE,” *Journal of Information Technology and Applications*, Volume 4, Issue 1, June 2014.
- [16] Jingjing Zhang and Nirwan Ansari, “On Assuring End-to-End QoE in Next Generation Networks: Challenges and a Possible Solution,” *IEEE Communications Magazine*, Volume 49, Issue 7, July 2011.
- [17] Ulrich Reiter et al., “Factors Influencing Quality of Experience,” In *Quality of Experience: Advanced Concepts, Applications and Methods*, 2014.
- [18] Sebastian Möller, M. Waltermann, and M. Garcia, “Features of Quality of Experience,” In *Quality of Experience: Advanced Concepts, Applications and Methods*, 2014.
- [19] Miguel Ríos Quintero and Alexander Raake, “Is taking into account the subjects’ degree of knowledge and expertise enough when rating quality?” In *Fourth International Workshop on Quality of Multimedia Experience (QoMEX)*, July 2012.
- [20] Telchemy, “Voice Quality Measurement,” In Voice over IP Performance Management, 2014. <http://www.telchemy.com/appnotes/TelchemyVoiceQualityMeasurement.pdf>
- [21] Ofcom, *On-Demand Services: Understanding Consumer Choices*, October 2012.
- [22] Andreas Sackl et al., “Wireless Vs. Wireline Shootout: How User Expectations Influence Quality of Experience,” In *Fourth International Workshop on Quality of Multimedia Experience (QoMEX)*, July 2012.
- [23] Florence Agboma and Antonio Liotta, “Addressing User Expectations in Mobile Content Delivery,” *Mobile Information Systems*, Volume 3, Issue 3–4, December 2007.
- [24] David S. Hands, “A Basic Multimedia Quality Model,” *IEEE Transactions on Multimedia*, Volume 6, No. 6, December 2004.
- [25] Stefan Winkler and Christof Faller, “Maximizing Audiovisual Quality at Low Bitrates,” *Proceedings of International Workshop on Video Processing and Quality Metrics for Consumer Electronics*, January 2005.

- [26] Tobias Hossfeld et al., “Initial Delay Vs. Interruptions: Between the Devil and the Deep Blue Sea,” In *Fourth International Workshop on Quality of Multimedia Experience (QoMEX)*, July 2012.
- [27] Sebastian Egger et al., “Waiting Times in Quality of Experience for Web Based Services,” In *Fourth International Workshop on Quality of Multimedia Experience (QoMEX)*, July 2012.
- [28] Peter Reichl, Patrick Maillé, Patrick Zwickl, and Andreas Sackl, “A Fixed-Point Model for QoE-Based Charging,” In *Proceedings of the 2013 ACM SIGCOMM Workshop on Future Human-Centric Multimedia Networking*, 2013.
- [29] Andrew Perkis, Peter Reichl, and Sergio Beker, “Business Perspectives on Quality of Experience,” In *Quality of Experience*, 97–108, 2014.
- [30] Florence Agboma, “Quality of Experience Management in Mobile Content Delivery Systems,” University of Essex, Ph.D. Thesis, 2009.
- [31] Martina Angela Sasse et al., “Remote Seminars through Multimedia Conferencing: Experiences from the Mice Project,” In *INET ’94*, 1994.
- [32] ITU-T, “Methods for the Subjective Assessment of Video Quality, Audio Quality and Audiovisual Quality of Internet Video and Distribution Quality Television in Any Environment,” Recommendation P.913, 2014.
- [33] ITU-R, “Methodology for the Subjective Assessment of the Quality of Television Pictures,” Recommendation BT.500, 2012.
- [34] ITU-T, “Subjective Video Quality Assessment Methods for Multimedia Applications,” Recommendation P.910, 2008.
- [35] ITU-R, “General Methods for the Subjective Assessment of Sound Quality,” Recommendation BS.1284-1, 2003.
- [36] Jesus Gutierrez et al., “Subjective Evaluation of Transmission Errors in IPTV and 3DTV,” In *IEEE Visual Communications and Image Processing (VCIP)*, 2011.
- [37] Nicolas Staelens et al., “Assessing Quality of Experience of IPTV and Video on Demand Services in Real-Life Environments,” *IEEE Transactions on Broadcasting*, Volume 56, Issue 4, December 2010.

- [38] Hendrik Knoche, John D. McCarthy, and Martina Angela Sasse, “Can Small Be Beautiful?: Assessing Image Resolution Requirements for Mobile TV,” In *ACM International Conference on Multimedia*, 2005.
- [39] Tobias Hossfeld et al., “Quality of Experience Management for YouTube: Clouds, Fog and the Aquareyoum,” *PIK – Praxis der Informationsverarbeitung und Kommunikation*, Volume 35, No. 3 2012.
- [40] ITU-R, “Parametric Non-Intrusive Assessment of Audiovisual Media Streaming Quality,” Recommendation P.1201, 2013.
- [41] ITU-T, “Objective Perceptual Multimedia Video Quality Measurement of HDTV for Digital Cable Television in the Presence of a Full Reference,” Recommendation J.341, 2011.
- [42] Marcus Barkowsky et al., “Hybrid video quality prediction: reviewing video quality measurement for widening application scope,” In *Multimedia Tools and Applications*, Volume 74, No. 2, Springer, 2015.
- [43] ITU-T, “Methods, Metrics and Procedures for Statistical Evaluation, Qualification and Comparison of Objective Quality Prediction Models,” Recommendation P.1401, 2012.
- [44] Kjell Brunnström, David Hands, Filippo Speranza, and Arthur Webster, “VQEG Validation and ITU Standardization of Objective Perceptual Video Quality Metrics [Standards in a Nutshell],” *IEEE Signal Processing Magazine*, Volume 26, No. 3, May 2009.
- [45] BBC News, “Xbox and PlayStation resuming service after attack,” <http://www.bbc.co.uk/news/uk-30602609>
- [46] Athula Balachandran, Vyas Sekar, Aditya Akella, Srinivasan Seshan, Ion Stoica, and Hui Zhang “Developing a Predictive Model of Quality of Experience for Internet Video,” In *ACM SIGCOMM Computer Communication Review*, Volume 43, No. 4, October 2013.
- [47] Florin Dobrian, Asad Awan, Dilip Joseph, Aditya Ganjam, Jibin Zhan, Vyas Sekar, Ion Stoica, and Hui Zhang, “Understanding the Impact of Video Quality on User Engagement,” *ACM SIGCOMM Computer Communication Review*, Volume 41, No. 4, August 2011.

- [48] ITU-T, “Mean Opinion Score Interpretation and Reporting,” Recommendation P.800.2, 2013.
- [49] Akira Takahashi et al., “Standardization activities in the ITU for a QoE assessment of IPTV,” *IEEE Communications Magazine*, Volume 46, No. 2, February 2008.
- [50] ITU-T, “Definitions of terms related to quality of service,” Recommendation E.800, September 2008.
- [51] ITU-T, “Vocabulary for performance and quality of service,” Recommendation P.10, July 2006.

WILLIAM STALLINGS is an independent consultant and author of numerous books on security, computer networking, and computer architecture. His latest book is *Foundations of Modern Networking: SDN, NFV, QoE, IoT, and Cloud* (Pearson, 2016). He maintains a resource site for computer science students and professionals at ComputerScienceStudent.com and is on the editorial board of *Cryptologia*. He has a Ph.D. in computer science from M.I.T. He can be reached at ws@shore.net

FLORENCE AGBOMA holds an M.Sc. and a Ph.D. in Electronic Systems Engineering from the University of Essex, United kingdom, in 2005 and 2009, respectively. Her Ph.D. research focused on Quality of Experience for Mobile Content Delivery Systems. She currently works as a Technology Analyst at British Sky Broadcasting (BSkyB), London. Her interests include subjective video quality assessment, psychophysical methods, and quality of experience management across devices and platforms. E-mail: fagbom@googlemail.com

IANA Stewardship Transition Moves to Final Phase

An historic proposal for the global community to assume stewardship of the *Internet Assigned Numbers Authority* (IANA) functions, produced after nearly two years of work by the global Internet community, has been delivered to the U.S. Government for its consideration. The proposal would remove U.S. Government oversight over a set of fundamental Internet administrative functions, including management of the global pool of Internet number resources (IPv4 and IPv6 addresses and *Autonomous System Numbers*), and replace it with a set of arrangements for community-based oversight.

The proposal, developed by the *IANA Stewardship Transition Coordination Group* (ICG), is based on input from three operational communities, including the *Internet Number Community* (those with an interest in the global management of Internet number resources). The contributions of the Internet Number Community were coordinated via a *Consolidated RIR IANA Stewardship Proposal* (CRISP) Team made up of community members drawn from each of the five RIR regions.

While the ICG published the final draft of its proposal in October 2015, elements of the proposal relied upon the adoption of a set of recommendations regarding the accountability of ICANN to its community. These recommendations were developed separately by a *Cross Community Working Group on Enhancing ICANN Accountability* (CCWG) and were adopted by the ICANN Board at its meeting in March 2016 in Marrakech, Morocco. The Board was at that point able to pass on both the ICG and CCWG documents to the *National Telecommunication and Information Administration* (NTIA), an agency of the U.S. Government.

The U.S. Government will now review the proposal to ensure that it meets the criteria set out by the NTIA when they first announced their intention, in March 2014, to pass stewardship of the IANA functions to the global community. If approved, the *Regional Internet Registries* (RIRs) and ICANN will continue their work towards implementation of the proposal, which will be completed prior to the expiration of ICANN's current contract with NTIA in September 2016.

For more information, see:

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

<https://www.nro.net/nro-and-internet-governance/iana-oversight/about-the-proposal>

<http://www.ianacg.org/icg-files/documents/IANA-transition-proposal-final.pdf>

<https://community.icann.org/display/acctcrosscomm/CCWG+on+Enhancing+ICANN+Accountability>

Call for Papers

The *Internet Protocol Journal* (IPJ) is a quarterly technical publication containing tutorial articles (“What is...?”) as well as implementation/operation articles (“How to...”). The journal provides articles about all aspects of Internet technology. IPJ is not intended to promote any specific products or services, but rather is intended to serve as an informational and educational resource for engineering professionals involved in the design, development, and operation of public and private internets and intranets. In addition to feature-length articles, IPJ contains technical updates, book reviews, announcements, opinion columns, and letters to the Editor. Topics include but are not limited to:

- Access and infrastructure technologies such as: Wi-Fi, Gigabit Ethernet, SONET, xDSL, cable, fiber optics, satellite, and mobile wireless.
- Transport and interconnection functions such as: switching, routing, tunneling, protocol transition, multicast, and performance.
- Network management, administration, and security issues, including: authentication, privacy, encryption, monitoring, firewalls, troubleshooting, and mapping.
- Value-added systems and services such as: Virtual Private Networks, resource location, caching, client/server systems, distributed systems, cloud computing, and quality of service.
- Application and end-user issues such as: E-mail, Web authoring, server technologies and systems, electronic commerce, and application management.
- Legal, policy, regulatory and governance topics such as: copyright, content control, content liability, settlement charges, resource allocation, and trademark disputes in the context of internetworking.

IPJ will pay a stipend of US\$1000 for published, feature-length articles. For further information regarding article submissions, please contact Ole J. Jacobsen, Editor and Publisher. Ole can be reached at ole@protocoljournal.org or olejacobsen@me.com

The Internet Protocol Journal is published under the “CC BY-NC-ND” Creative Commons Licence. Quotation with attribution encouraged.

This publication is distributed on an “as-is” basis, without warranty of any kind either express or implied, including but not limited to the implied warranties of merchantability, fitness for a particular purpose, or non-infringement. This publication could contain technical inaccuracies or typographical errors. Later issues may modify or update information provided in this issue. Neither the publisher nor any contributor shall have any liability to any person for any loss or damage caused directly or indirectly by the information contained herein.

Supporters and Sponsors

Publication of this journal is made possible by:

Supporters



Diamond Sponsors



Ruby Sponsor



Sapphire Sponsors



Emerald Sponsors



Corporate Subscriptions



Individual Sponsors

Lyman Chapin, Steve Corbató, Dave Crocker, Jay Etchings, Martin Hannigan, Hagen Hultzs, Dennis Jennings, Jim Johnston, Merike Kaeo, Bobby Krupczak, Richard Lamb, Tracy LaQuey Parker, Bill Manning, Andrea Montefusco, Tariq Mustafa, Mike O'Connor, Tim Pozar, George Sadowsky, Scott Seifel, Helge Skrivervik, Rob Thomas, Tom Vest, Rick Wesson.

For more information about sponsorship, please contact sponsor@protocoljournal.org

The Internet Protocol Journal
NMS
535 Brennan Street
San Jose, CA 95131

ADDRESS SERVICE REQUESTED

The Internet Protocol Journal

Ole J. Jacobsen, Editor and Publisher

Editorial Advisory Board

Fred Baker, Cisco Fellow
Cisco Systems, Inc.

Dr. Vint Cerf, VP and Chief Internet Evangelist
Google Inc, USA

Dr. Steve Crocker, Chairman
Internet Corporation for Assigned Names and Numbers

Dr. Jon Crowcroft, Marconi Professor of Communications Systems
University of Cambridge, England

Geoff Huston, Chief Scientist
Asia Pacific Network Information Centre, Australia

Olaf Kolkman, Chief Internet Technology Officer
The Internet Society

Dr. Jun Murai, Founder, WIDE Project, Dean and Professor
Faculty of Environmental and Information Studies,
Keio University, Japan

Pindar Wong, Chairman and President
Verifi Limited, Hong Kong

The Internet Protocol Journal is published quarterly and supported by the Internet Society and other organizations and individuals around the world dedicated to the design, growth, evolution, and operation of the global Internet and private networks built on the Internet Protocol.

Email: ipj@protocoljournal.org
Web: www.protocoljournal.org

The title "The Internet Protocol Journal" is a trademark of Cisco Systems, Inc. and/or its affiliates ("Cisco"), used under license. All other trademarks mentioned in this document or website are the property of their respective owners.

Printed in the USA on recycled paper.

