5G Technology

ICANN Office of the Chief Technology Officer

Alain Durand
OCTO-004
23 January 2020

TABLE OF CONTENTS

This document is part of the OCTO document series. Please see https://www.icann.org/resources/pages/octo-publications-2019-05-24-en for a list of documents in the series. If you have questions or suggestions on any of these documents, please send them to octo@icann.org.

1 Executive Summary

The Internet is increasingly mobile. The next wave of devices Internet-connected devices won't be limited to computers or smartphones with an actual human being in front of the screen. It will include Machine-to-Machine communication with Internet of Things (IoT) devices. With those two trends in mind, the introduction of 5G (the fifth generation wireless technology for cellular networks) is of special importance to the Internet in general and ICANN in particular.

The fundamental question being asked by the introduction of 5G boils down to: Is the current model of the Internet (known as the TCP/IP protocol suite) still relevant in a 5G world? This question for ICANN translates into the two following questions:

- Is DNS still operable in a 5G world, especially considering latency sensitive applications?
- Is the set of unique identifiers ICANN helps coordinate still useful, or does 5G bring with it the need for a different set of identifiers?

Over the years, there have been a number of academic research projects aimed at redefining networking for a "future internet". More importantly, any new attempt to replace TCP/IP should probably consider decades-long timelines and any alleged benefits would have to outweigh the complexities and cost of such transition. Furthermore, ICANN notes that the IETF has already done extensive work to make IP work in constrained environments, such battery-operated devices or very low power/very low bandwidth networks in the 6lopan and successor 6lo working groups. Other IETF efforts, such as those done in the QUIC working group, evolve the transport layer protocols to provide, among other things, stream multiplexing and low-latency connection establishment.

There does not appear to be a clear need for a new identifier system for classic user-oriented applications using 5G. However, IoT is a domain that could benefit from new global identifiers, especially some that could better handle privacy. Such identifiers can be implemented directly within the DNS.

DNS resolution latency and caching are a network operation/optimization concern, not an architectural issue. ICANN has the following recommendations for 5G networks:

- DNS caches for latency-sensitive 5G apps should be as local as possible and have aggressive prefetching configured.
- A distributed caching system might help maximize the efficiency of the overall DNS resolution system.
- IoT application developers looking to minimize the effect of DNS latency may want to investigate adapting their applications to query DNS data well ahead of connection establishment.

ICANN believes the model of a single Internet, based on a global system of unique identifiers, is the best way to maximize the benefits the Internet can bring. There is a risk that popular platforms could evolve to leverage 5G Network Slicing using their own, private, identifier system. If that were to happen, the Internet would fracture and only the long tail of lesser-known applications will keep using the Internet global system of unique identifiers.

2 Introduction

This memo will take a look at 5G (the fifth generation wireless technology for cellular networks) from a technical perspective, asking the questions: what does 5G change, if anything, to the Internet architecture and protocols such as TCP/IP? What would be the impact on the system of unique identifiers that ICANN helps coordinate, notably the Domain Name System (DNS)?

3 Latency: Is the DNS workable in a 5G low latency environment?

3.1 Background

The discussions related to mobile architectures have frequently been framed by a balance between operators and vendors. The operators drive large parts of the requirements, whereas the vendors create the suitable technology to match those requirements. Operators are interested in bringing new entrants as potential suppliers into the market and existing vendors are interested in keeping (and growing) their market share. Each new generation of mobile communications technology brings a new architecture (or evolution of existing architecture) with the promise of new services and business opportunities. These new technologies are touted as an avenue for new entrants (vendors, operators or third parties) to disrupt the market.

In 5G, incumbent vendors have initially pushed to maintain a centralized architecture while improving the radio. New entrants have pushed, since the conceptualization of 5G, for an edge computing architecture, promoting Software Defined Network (SDN) and Network Function Virtualization (NFV). (Edge computing is a design to decrease bandwidth and delay by moving needed resources closer to the systems requesting it.) The background of this contention is the promise of lower capital expenditure and the possibility to offer new services, potentially delivered a-la-carte, that could be launched at reduced cost. Among those new services would be the possibility to offer Ultra Reliable Low Latency Communications (URLLC), i.e., sub-5ms or sub-10ms, for self-driving cars, Vehicle to Everything (V2X), and Augmented Reality (AR)/Virtual Reality (VR) applications.

This move toward NFV started in 2012, before 5G, when the European Telecommunications Standards Institute (ETSI) created the NFV Industry Specification Group (ISG). The NFV technology has somewhat matured since then and incumbents now also offer a large part of their product portfolio as Virtualized Network Functions (VNFs). Edge clouds, under various definitions, came from the operator community. There have been different edge cloud initiatives both in standardization, e.g. ETSI Multi-Access Edge Computing (MEC), and in open source communities where, for example, the Linux Foundation Edge Foundation (LF Edge)[1] provides an umbrella to “establish an open, interoperable framework for edge computing independent of hardware, silicon, cloud or operating system”.

[1] See https://www.lfedge.org

A number of articles[2,3] have been published recently by some of the new vendors, pushing the idea that traditional DNS is not compatible with 5G extra low latency applications, i.e., URLLC. The key argument is that a round trip time to a "regional" data center adds unacceptable extra latency. The solution those vendors push is to put DNS resolvers in the edge computing platforms instead of centralizing the DNS resolution to a national or regional data center.

3.2 Discussion

The 5G radio link is just the first leg of any communication. One of the goals of 5G is to reduce radio latency to under 5ms or 10ms Round Trip Time (RTT). The communication is then transported (backhauled) from a radio tower to a data center. Assuming a regional data center 1,000km away from an edge facility, connected over fiber, an additional 10ms RTT would be added. However, to this, the time the resolver will take to process the DNS queries would need to be taken into consideration, with a similar calculation time also needing to be added to the edge resolvers. If the destination of the TCP or UDP connection following the DNS resolution is a host situated outside of the vicinity of the edge data center, this extra 10ms delay, only happening once at connection establishment, would likely be negligible. However, if the destination host is located within the edge data center (e.g., inter-VNF communications) or connected to it via 5G (e.g., self-driving trucks in a mine), this delay may be significant. In such cases, hosting the DNS resolvers in the edge data center might make sense. Another option is for the application to prefetch DNS data at boot time to reduce any delay at connection time. This is doable in an industrial environment where the set of destinations a device will connect to is known ahead of time. For regular Internet connections, DNS latency does not appear to be an issue.

Placing general purpose DNS resolvers closer to the users in an edge data center would have the effect of reducing the hit rate on the resolver DNS cache. This impact could be mitigated by installing a distributed caching system, such as a hierarchy of regional and national caches or deploying a form of DNS record prefetch[4,5]. However, there is another effect of 5G mobility on DNS caching that needs to be taken into consideration. If a mobile is moving, it might need to be redirected from time to time to a different, closer edge datacenter in order to maintain sub-10ms latency. This redirection might be taking place via a call to the DNS made by a mobility-aware application. In such a case, the DNS response will be calculated from the new geographical position of the mobile device. This is a variation of DNS-based load balancing practiced by many Content Delivery Networks (CDN) today, with the difference that the response should not be cached by the mobile device. This is standard DNS engineering practice, where the Time-To-Live (TTL) of certain DNS records is set to 0.

3.3 ICANN position

2 See https://www.open-xchange.com/about-ox/ox-blog/article/dns-latency-in-a-5g-network/
3 See https://www.infoblox.com/wp-content/uploads/infoblox-solution-note-infoblox-dns-for-5g.pdf
4 An example of DNS prefetch can be found at https://www.researchgate.net/publication/270571591_PREFETCHing_to_optimize_DNSSEC_deployment_over_large_Resolving_Platforms
5 Prefetching is already implemented in many resolver implementations.

DNS resolution latency and caching are a network operation/optimization concern, not an architectural issue. As such, ICANN believes that the DNS is workable in a 5G low-latency environment. ICANN has the following recommendations:

- DNS caches for latency-sensitive 5G apps should be as local as possible and have aggressive prefetching configured.
- A distributed caching system might help maximize the efficiency of the overall DNS resolution system.
- IoT application developers looking to minimize the effect of DNS latency may want to investigate adapting their applications to query DNS data well ahead of connection establishment.

4 Network Slicing: Is there a risk of fragmentation of the Internet's unique identifier system?

4.1 Background

5G promotes the concept of Network Slicing to abstract network resources and network functions. It enables carriers to build a single physical network that can account for very different use cases: high bandwidth applications (e.g., streaming), low bandwidth applications (e.g., connecting Internet of Things (IoT) devices) with low latency requirements, enterprise extranets, etc. Network Slicing is a new term for an old concept. From 2G onwards, mobile networks have had capabilities called "Packet Data Protocol (PDP) Contexts/Packet Data Network (PDN) Connections" that are selected based on an Access Point Name (APN). APNs and their underlying PDP/PDN infrastructures have been used for enterprise customers providing direct connectivity for their internal networks. 5G goes further, allowing Quality of Server (QoS) parameters to be set. Operators would be capable of setting aside bandwidth based on customer QoS requirements to provide a network slice running on top of a single physical infrastructure. Currently, one of the main use cases for Network Slicing is "industry 4.0," a term used to describe scenarios in which an operator can practically offer an "own network" for a factory or another industry where they have a guaranteed bandwidth and, especially, deterministic latency. This functionality would allow the various industries to move from proprietary wireline infrastructure to a more flexible wireless network technology.

4.2 Discussion

5G is defined in 3GPP. 3GPP Service and System Aspect, Architecture Working Group (SA2) has defined Network Slice Instances (NSIs), where each NSI contains several Network Slice Subnet Instances (NSSIs). The 3GPP System Architecture for 5G Systems Technical Specification (TS 23.501) defines the Network Slice Selection Assistance Information (NSSAI), which is used to assist the User Equipment (UE) in the Network Slice selection and the Service Slice Type (SST). This standardization of network slices is still in its early phases. It is ready for statically provisioned network slices, but more work is needed to enable dynamically provisioned network slices in an SDN-type approach.

Provisioning dedicated bandwidth for a specific network is not neutral to the operator. It takes away resources available for the common pool. It is based on the notion that the revenues generated by the spectrum set aside for the targeted customers will more than offset the loss of revenue generated by the corresponding drop in available spectrum for generic customers. How Network Slicing will actually be implemented and priced by cellular operators is still unclear. Some of the technical and business challenges are reminiscent of Constant Bit Rate (CBR)/Available Bit Rates (ABR)/Variable Bit rate (VBR) offers on ATM networks in the late 1990s. Operators back then were interested in offering such QoS services but were reluctant to let their customers dynamically provision those for fear of over-provisioning the network.

Going beyond "industry 4.0," Network Slicing could also be used to segregate between multiple "services/applications." There is a possibility that a combination of specialized applications plus 5G Network Slicing and classic Virtual Private Network (VPN)/Virtual Routing and Forwarding (VRF) technologies could be deployed to create large extranets that would connect users independently of the common Internet to popular well-known services such as Facebook, Netflix, Amazon, and others. That is contrary to today, where users can gain access to all of these services through a single network. Instead, a user application would get access to the "Facebook" slice, the "Netflix" slice, or the "Amazon" slice to get better service when accessing these services. This could be an evolution of the current model where those over-the-top players already deploy CDN caches close to the eyeballs in Internet Service Provider (ISP) networks. Network slices deployed this way would offer a dedicated network connecting with constrained QoS parameters, i.e., no longer "best effort" connectivity, the handset connected directly to the over-the-top player network. In other words, the handset would no longer connect to Facebook, Netflix, or Amazon over the Internet but would be part of those respective networks.

On top of the net neutrality aspects of such deployment, the multiplication of such per-application network slices would be a radical departure from a key concept of the Internet: one network with multiple applications. In such a model, there would be multiple dedicated networks, one per application. Slices could use names and addresses coming from the globally unique identifier system ICANN helps to coordinate, but this is not a technical requirement. At the request of the application owner, such slices could be deployed using the owner's dedicated set of identifiers, address space, and name space. This scenario would further increase the fragmentation of the Internet.

There is no indication that this scenario is planned in the initial or subsequent 5G rollout plans. Further, operators may choose to rollout network slices using global unique identifiers. As such, the risk created by Network Slicing on the fragmentation of the Internet appears low at the moment. For the scenario described above to happen, a content provider such as Facebook would have to convince 5G operators that represent a significant percentage of Facebook's customer base to create a network slice to Facebook's requirements and then connect that slice to their private content delivery network. The likelihood of such a scenario probably depends on the overall balance of power between ISPs and over the top players. Back in the late 2000s, a similar situation existed. Content providers wanted to deploy cache engines deep inside ISP networks. Dedicated bandwidth in the ISPs had to be reserved to feed the cache directly from the content providers. The question was: who pays for that bandwidth? The content provider who benefits from eyeballs closer to its content, or the ISP who benefits from content closer to its eyeballs? A combination of a drop in long-distance bandwidth cost and a rise in power of those content providers means that such caches are now a reality, commonly deployed deep inside ISP networks. Creating private network slices might just be a repeat of that discussion.

4.3 ICANN position

ICANN believes the model of a single Internet, based on a global system of unique identifiers, is the best way to maximize the benefits the Internet can bring. There is a risk that popular platforms could evolve to leverage Network Slicing using their own identifier system. If that were to happen, the Internet would fracture and only the long tail of lesser-known applications will keep using the Internet global system of unique identifiers.

5 Will phone numbers still be relevant with 5G? Will 5G result in the introduction of new sets of identifiers? Will those identifiers be in the DNS?

5.1 Background

The reliance on Voice over LTE (VoLTE) for basic voice service, coupled with the now dominant roles of services such as WhatsApp, Telegram, Facetime, and others might suggest that phone numbers are a relic of the past.

IoT communications may also require very different sets of identifiers, either ephemeral or persistent, associated with various privacy and security requirements.

The question is: in 5G, what new sets of identifiers, if any, are required; would those identifiers be based on the DNS or not; and is IP(v4 or v6) still relevant?

5.2 Discussion

E.164[6] numbers are used within cellular networks only to identify end-user devices. Internally, since 2G, cellular networks have used another identifier, the International Mobile Subscriber Identity (IMSI), to route calls. Similarly, WhatsApp and other similar applications make use of an E.164 telephone number to identify a user but use IP to move data and place calls. As such, E.164 lives on as end-user identity.

Enum[7] has not been deployed much outside of telephone number portability.

There are no bridges between the various Instant Messaging (IM) systems and social media platforms. The reason being that those platforms compete with each other and see no value in interoperability. For example, if a user on WhatsApp wanted to communicate with a user on Telegram, at least one of them would have to sign-up for the other service and download the

[6] E.164 is an ITU-T recommendation that defines an international numbering plan for the world-wide Public Switched Telephone Network (PSTN).
[7] Enum is a mapping of a E.164 telephone number into a URI through the DNS. Enum is defined in https://tools.ietf.org/html/rfc6116 and https://tools.ietf.org/html/rfc6117.

appropriate app. As such, a generic directory system that would introduce a new set of identifiers that maps to a specific IM or social media platform would be of little use.

Today, IoT device manufacturers typically use their own proprietary systems to identify and address those devices. The manufacturers have many schemes to choose from: the serial number of the device, the IMEI[8] number, a MAC address, a DOA identifier, or something entirely proprietary. Most of those identifiers are tied to the hardware and are essentially persistent. This persistence may cause privacy concerns if the mapping of the persistent identifier to the owner/user of the device can be obtained. To address this concern, a new set of privacy-aware, ephemeral identifiers might be required. The usage of such ephemeral or persistent identifiers hosted in the DNS has been studied by ICANN Office of the CTO, and prototypes have been developed and presented at the ICANN meeting in Abu Dhabi in November 2017 to demonstrate the feasibility of using DNS for IoT identifiers.

5G, just like 4G and the previous iterations before, heavily leverages IPv4, IPv6 addresses, and domain names. The benefits of the introduction of any new identifier system would need to outweigh the complexities and the costs of developing and deploying such a new system while maintaining interoperability with the existing domain names/IP addresses.

5.3 ICANN position

There does not appear to be a clear need for a new identifier system for classic user-oriented applications using 5G. However, IoT is a domain that could benefit from new global identifiers, especially some that could better handle privacy. Such identifiers can be implemented directly within the DNS.

6 ITU-T Focus Group Network 2030

6.1 Background

While not directly linked to 5G (yet), ITU-T has started a new effort in the Network 2030 Focus Group[9]. The stated goal is to define a new Layer 3 network protocol (a replacement for IP). A white paper[10] and a technical report[11] were published in 2019 with a starting point that TCP/IP is not suitable for future applications such as holographic communications and machine-to-machine communications. A key element highlighted by Network 2030 is an access control mechanism to move beyond best effort and guarantee delay and jitter. An Application Programming Interface (API) is called for, to enable applications to program the network directly before starting a communication, as opposed to measuring network propagation characteristics and adapting to it. Another design element is to group communication flows so they can share the same fate in case of congestion. A further one is to allow the network layer elements to 'downgrade' some streams of traffic in case of congestion.

[8] The IMSI (International Mobile Subscriber Identity) is a code used by the cellular operator to identify the SIM on the mobile network. The IMEI (International Mobile Station Equipment Identity) is an international "Serial number" for the device itself.
[9] See https://www.itu.int/en/ITU-T/focusgroups/net2030/Pages/default.aspx
[10] See https://www.itu.int/en/ITU-T/focusgroups/net2030/Documents/White_Paper.pdf
[11] See https://www.itu.int/en/ITU-T/focusgroups/net2030/Documents/Deliverable_NET2030.pdf

Instead of targeting a ubiquitous network, this Network 2030 effort is aimed at creating tailored sets of networks for specific verticals. A normal situation for a network device could be to be connected to “multiple” specialized Internets, instead of just one.

Note: Network 2030 is aiming wired line infrastructure with an eye to wireless networks in “Beyond 5G”/6G architectures.

6.2 Discussion

Claims that TCP/IP is not suitable for some types of new-and-coming applications are not new. As a matter of fact, they resurface each time a new access network technology, such as DSL, Fiber-To-The-Home (FTTH), 3G, 4G, 5G, etc. comes along. The track-record of efforts to improve on TCP transport protocol shows that most efforts eventually realize that TCP is still the best approach. However, there may be a future where this is no longer the case, potentially when managing connections with spacecrafts, planets, and other far-flung objects. This does not mean that there are no needs for new transport protocol other than TCP; the IETF is well underway in its effort to standardize the QUIC[12] transport protocol that provides, among other things, stream multiplexing and low-latency connection establishment.

Discussions about the necessity (or not) of admission control mechanisms to guarantee QoS have been going on since the dawn of networking. In the last few decades, the answer to those questions has simply been “more bandwidth,” rather than going back to a connection-oriented networking model such as the legacy telephony network, as advocated by Network 2030.

Grouping streams together and providing a new API to allow an application to better communicate requirements to the underlying network does not require a new layer protocol. Many efforts have been started in those directions in the IETF. Also worth noting is how stream bandwidth adaptation has been implemented in video content delivery networks for many years, using application layer relays.

A large part of the focus of the Network 2030 documents are machine-to-machine communications that require sub 10ms or even sub 1ms round-trip time (RTT). As the technical paper published by the ITU-T FG2030 points out, this requirement is gated by the speed of light. 10ms RTT is roughly 1,000km, 1ms is 100km. As such, we are here talking about local (or at best regional) area networks where specialized technologies and engineering practices could be deployed to address the specific requirements without impacting the global Internet.

The idea of devices connected to “multiple” specialized Internets share some of the same potential issues discussed earlier in this document.

It should be noted that the requirements and use cases presented in Network 2030 are not very detailed and do not lay a very solid technical foundation to make the case that a new networking protocol suite is actually needed. As such, this work might be considered premature and edging more on science fiction (e.g. holographic communications) than based on current and actual networking issues. It should also be noted that this focus group does not appear to be representative of a cross-section of the entire industry.

[12] See https://datatracker.ietf.org/doc/draft-ietf-quic-transport/

The Focus Group 2030 was supposed to finish its work in November 2019 but was granted a one-year extension.

One thing to keep in mind is how hard the transition from IPv4 to IPv6 has been. More than 20 years since IPv6 was first standardized, IPv6 is still very far from being universally deployed. Any new attempt to replace TCP/IP should probably consider decades-long timelines.

6.3 ICANN position

Over the years, there have been a number of academic research projects aimed at redefining networking for a "future internet". However, the development and standardization of network and transport layer protocols such as the TCP/IP protocol suite has traditionally fit in the realm of the IETF, not the ITU-T. More importantly, any new attempt to replace TCP/IP should probably consider decades-long timelines and any alleged benefits would have to outweigh the complexities and cost of such transition.

7 Are there opportunities for non-IP solutions at the edge in 5G?

7.1 Background

Non-IP solutions at the edge of 5G have been proposed to address perceived latency or alleged limitations in the IP model to support constrained environments such as battery-operated devices or low energy/low bandwidth networks, or properly support latency sensitive applications.

7.2 Discussion: Can non-IP solutions be deployed in 5G?

This question can be broken down in different ways:

Can two 5G devices, connected to the same edge, communicate directly, possibly using non-IP solutions at layer 3?

Yes, today. 3GPP release 15[13] has defined an Ethernet Packet Data Unit (PDU), so two devices connected this way could either talk to each other directly at L2 over Ethernet or implement any layer 3 protocol of their choice, not necessarily IP. Such devices would have to implement a specialized protocol stack. This is possible in a vertical market such as Machine to Machine (M2M) communications in an Industry 4.0 context.

Can two 5G devices, connected to the same edge, communicate directly, possibly using some new layer 2 extensions?

[13] See https://www.3gpp.org/release-15

Possibly, in the near future. 3GPP is looking at defining a profile to support Time Sensitive Networking (TSN) extensions[14] to the IEEE 802.1[15] standard in 3GPP release 16[16]. As in the previous case, devices would have to implement a specialized protocol stack. This is possible in a vertical market such as Machine-to-Machine (M2M) communications in an Industry 4.0 context.

Can a 5G device use non-IP technology to communicate with a server within the edge data center?

Yes. The server would have to implement a specialized protocol stack.

It should be noted that the above cases only really apply for private communications between devices under the same (or related) administrative control, and relatively close geographically. If the two endpoints are far away, any perceived latency benefits of replacing IP disappears because of the speed of light limitation. If the communication involves entities under different administrative controls, the complexities of setting up the technical connections and the proper business relationships between the different entities would make such scenario difficult.

7.3 Discussion: Can IP work on constrained devices?

The Internet Engineering Task Force (IETF) has been very active in making IP work on constrained networks. In particular, the working group 6lowpan[17] and its successor 6lo[18] have defined extensions to enable IP on resource constrained devices, such as battery-operated ones or devices using very low bandwidth radio.

Among supported link layer technologies, we can mention: IEEE 802.15.4[19] supported in RFC4944[20], ITU-T G.9959[21] (Zwave) supported in RFC7428[22], Bluetooth Low Energy[23] (BLE) supported in RFC7668[24], Digital Enhancement Cordless Telecommunications/Ultra Low Energy[25] (DECT-ULE) supported in RFC8105[26], Master Slave Token Passing[27] (MS/TP) supported in RFC8163[28], Near Field Communications[29] (NFC) supported in draft-ietf-6lo-nfc[30],

14 See https://1.ieee802.org/tsn/
15 See https://1.ieee802.org
16 See https://www.3gpp.org/release-16
17 See https://datatracker.ietf.org/wg/6lowpan/about/
18 See https://datatracker.ietf.org/wg/6lo/about/
19 See http://www.ieee802.org/15/pub/TG4.html
20 See https://tools.ietf.org/html/rfc4944
21 See https://www.itu.int/rec/T-REC-G.9959
22 See https://tools.ietf.org/html/rfc7428
23 See https://www.bluetooth.com/learn-about-bluetooth/bluetooth-technology/radio-versions/
24 See https://tools.ietf.org/html/rfc7668
25 See https://www.ulealliance.org
26 See https://tools.ietf.org/html/rfc8105
27 ANSI standard 135-2016, BACNET, a data communication protocol for building automation and control networks
28 See https://tools.ietf.org/html/rfc8163
29 See https://www.iso.org/standard/56692.html
30 See https://datatracker.ietf.org/doc/draft-ietf-6lo-nfc/

Power Line Communication[31] (PLC) supported in draft-ietf-6lo-plc[32]. Various techniques are applied, ranging from header compression RFC3095[33], RFC6282[34], RFC7400[35], link layer fragmentation and reassembly via an adaptation layer RFC4944[36], protocol optimization (e.g. IPv6 neighbor discovery optimization RFC6775[37], routing optimization in constrained networks, RFC6550[38]). Use cases describing IP in those constrained environments are described in the IETF document draft-ietf-6lo-use-cases[39].

7.4 Discussion: How to take into account latency sensitive applications in TCP/IP?

The Institute of Electrical and Electronics Engineers (IEEE) is defining Time Sensitive Networking (TSN) extensions[40] to the IEEE 802.1[41] standard. The IETF, in collaboration with IEEE 802.1, has chartered the Deterministic Networking (DETNET) working group. This working group's charter is to work on "deterministic data paths that operate over Layer2 bridged and Layer3 routed segments, where such path can provide bounds on latency, loss and packet delay variation (jitter), and high reliability".

Another angle of the IETF's work in that domain is the QUIC[42] transport protocol that provides, among other things, stream multiplexing and low-latency connection establishment.

7.5 ICANN position

Specific verticals using private 5G networks or network slices of public 5G networks are just a special case of proprietary networks. Within private, proprietary networks, people are free to leverage specific non-IP based technologies without any impact on the global Internet.

ICANN notes that the IETF has already done extensive work to make IP work in constrained environments, such battery-operated devices or very low power/very low bandwidth networks in the 6lopan and successor 6lo working groups, and to support latency sensitive applications in the DETNET working group in collaboration with the IEEE Time Sensitive Networking group. Another example of the IETF involvement in latency sensitive environment is the QUIC transport protocol that provides, among other things, stream multiplexing and low-latency connection establishment.

31 See https://standards.ieee.org/standard/1901-2010.html
32 See https://tools.ietf.org/html/draft-ietf-6lo-plc-01
33 See https://tools.ietf.org/html/rfc3095
34 See https://tools.ietf.org/html/rfc6282
35 See https://tools.ietf.org/html/rfc7400
36 See https://tools.ietf.org/html/rfc4944
37 See https://tools.ietf.org/html/rfc6775
38 See https://tools.ietf.org/html/rfc6550
39 See https://datatracker.ietf.org/doc/draft-ietf-6lo-use-cases/
40 See https://1.ieee802.org/tsn/
41 See https://1.ieee802.org
42 See https://datatracker.ietf.org/doc/draft-ietf-quic-transport/