

# The Internet Protocol Journal

September 2023

Volume 26, Number 2

A Quarterly Technical Publication for  
Internet and Intranet Professionals

## FROM THE EDITOR

### In This Issue

From the Editor .....	1
Introduction to 5G .....	2
LEO Satellites for Internet Access.....	31
Letter to the Editor.....	44
Fragments.....	45
Thank You! .....	48
Call for Papers.....	50
Supporters and Sponsors .....	51

Technologies used for accessing the Internet have evolved a great deal since my very first encounter with the ARPANET in 1976. Using a 110-baud Teletype machine, I accessed a computer at SRI International in Menlo Park, California, from my laboratory at the *Norwegian Defence Research Establishment* (NDRE) at Kjeller, Norway. Today, my Internet service is delivered by fiber-optic cable at 1 Gbps. Numerous other technologies for Internet access have emerged, and in this issue we look at two of them, namely 5G mobile systems and *Low Earth Orbit* (LEO) satellites.

In simple terms 5G can be described as a new set of cellular radio frequencies to allow for much faster data connections for mobile devices. Beyond this simplified explanation lies numerous details that are described in a two-part article by William Stallings. Part One introduces the standards and specifications that define 5G and describes the usage scenarios that 5G supports. Part Two, to be published in a future issue, will provide an overview of the structure and function of 5G networks. A third article on *Network Slicing*, which is closely related to 5G, will also be published in a future edition of this journal.

The world's first communication satellite, *Telstar 1*, which was launched in July 1962, provided proof-of-concept for both live television transmission and telephone service. Since that time, satellites have been deployed for many services, ranging from weather observations, navigation systems, and more recently Internet access. LEO satellites are particularly well-suited for Internet access because they offer coverage to remote areas without introducing substantial propagation delays as compared to other alternatives. Our second article, by Dan York and Geoff Huston, provides an overview of LEO systems for Internet access.

This journal now has around 1,000 print subscribers and just over 18,000 online subscribers who download their copy from our website. Given this shift in subscriber demographics, we will no longer be printing those long and cumbersome URLs in the references section at the end of each article. Instead, you can simply click on the references themselves using the PDF copy.

As always, we welcome your feedback and suggestions on anything you read in this journal. Letters to the Editor may be edited for clarity and length and can be sent to [ipj@protocoljournal.org](mailto:ipj@protocoljournal.org)

—Ole J. Jacobsen, Editor and Publisher  
[ole@protocoljournal.org](mailto:ole@protocoljournal.org)

You can download IPJ  
back issues and find  
subscription information at:  
[www.protocoljournal.org](http://www.protocoljournal.org)

ISSN 1944-1134

# Introduction to 5G

## Part One: Standards, Specifications, and Usage Scenarios

by William Stallings

5G is the fifth-generation technology for wireless cellular networks. A significant technological leap beyond the capabilities of the 4G networks that currently dominate available cellular network services, 5G delivers a substantial increase in peak and average speeds and capacity. A significant increase in download and upload speeds will enhance many existing use cases, including cloud-based storage, augmented reality, and artificial intelligence. It also will enable cell sites to communicate with a greater number of devices. Reduced latency enables edge computing and will transform *Internet of Things* (IoT) capabilities and application breadth.

This two-part article provides an introduction to 5G. Part One introduces the standards and specifications that define 5G and describes the usage scenarios that 5G supports. Part Two, to be published in a future issue, will provide an overview of the structure and function of 5G networks.

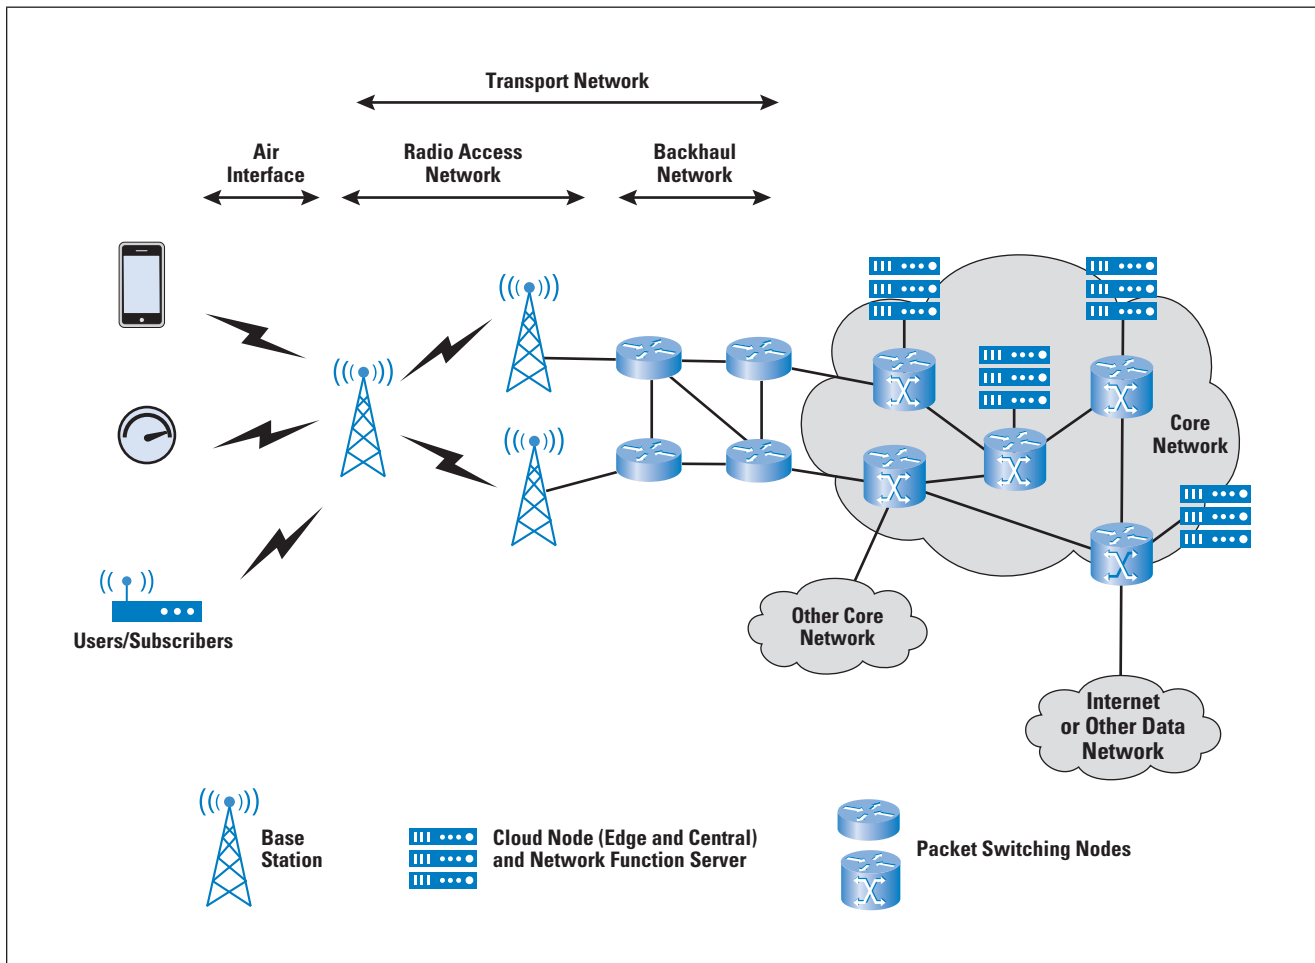
### Simplified View of a 5G Network

Figure 1 shows a very simplified view of a 5G cellular network; it provides a useful framework for discussing standards and specifications for the technology. Users, or subscribers, to the network are fixed or mobile wireless devices, referred to in 5G documents as *User Equipment* (UE). Examples of fixed wireless UE include a modem that serves as an access point from a home or office Wi-Fi or Ethernet network to the 5G network, and IoT devices such as sensors or surveillance cameras. Mobile UE includes cell phones, laptops, and other mobile or portable devices equipped with 5G capability.

The *Air Interface*, also called a *Radio Interface*, is the wireless link between UE and the nearest cellular base station. The air interface specifies the method for transmitting information over the air between base stations and mobile units, including protocols, frequency range, channel bandwidth, channel coding, and the modulation scheme.

The *Radio Access Network* (RAN) consists of a collection of base stations that provide service to the UE in a geographic region. The base station provides radio transmission and reception in one or more cells to or from the user equipment. A base station can have an integrated antenna or can be connected to an antenna by feeder cables. Each base station communicates with nearby base stations, generally wirelessly, to enable handoff of UE from one base station to another as the UE moves. The RAN also includes other management and transmission elements.

Figure 1: Simplified View of 5G Network



The *Core Network* is a backbone network that provides interconnection service between RANs in different regions; it provides access for UE to the Internet or other data networks and UE on other RANs. In addition, the core network implements numerous network functions that support user- and control-plane traffic and provides for *Quality of Service* (QoS) and management and orchestration. Core networks also generally provide both edge and central cloud services for 5G users.

The *Transport Network* is the collection of communication links that interconnect nodes of the RAN, as well as communication links connecting RAN elements to the 5G core network. Links between RAN elements and UE are generally not considered part of the transport network.

### Standards and Specifications for 5G

Many of the important developments in information and technology and communications, such as the Internet, IoT, Cloud Computing, and Virtualization, have been driven in part by international standards.

However, in all of these cases, much of the technology was developed and deployed in advance of universally agreed-upon standards. The case of 5G is quite different. Although a reasonably complete set of standards based on fixed specification is only just coming to fruition, the implementations and deployments that preceded these standards and specifications anticipated their final form. Throughout the 5G ecosystem, which includes device and component manufacturers, cellular network providers, network software providers, and application developers, the work done prior to the introduction of the first set of standards in 2020 closely follows what has ultimately been standardized. Going forward, there is universal agreement that 5G-related implementations will follow the standards.

Because an understanding of 5G depends on an understanding of the process by which the standards are developed and the content of those standards, the first part of this article provides an overview. It covers the two organizations that are responsible for the development of 5G: the *International Telecommunication Union* (ITU) and the *3rd Generation Partnership Project* (3GPP). In essence, the process of standards development for 5G follows this sequence:

1. The ITU has issued—and continues to issue—standards, called *Recommendations*, and other documents, called *Reports*, that define the overall concept for 5G, as well as the technical, performance, and service requirements for 5G.
2. Based on the ITU requirements, as well as requirements generated by national and regional standards organizations and market-based organizations, 3GPP has developed—and continues to develop—a detailed set of technical specifications for the implementation of 5G.
3. The ITU has translated these specifications into international standards (Recommendations) that dictate how 5G is being implemented.

This process is ongoing as further refinements and capabilities are added to the requirements and the technical specifications.

With respect to 5G, the two relevant components of ITU are the *ITU Radiocommunication* (ITU-R) Sector and the *ITU Telecommunication Standardization* (ITU-T) Sector. In general terms, ITU-R issues standards related to user requirements and the air interface. ITU-T issues standards related to the RAN, the transport network, and the core network.

#### **ITU-R and IMT-2020**

Perhaps the most prominent initiative by ITU-R is the *International Mobile Telecommunications* (IMT) project. IMT is the generic term the ITU community uses to designate broadband mobile systems. It encompasses *IMT-2000*, *IMT-Advanced*, and *IMT-2020* collectively, which correspond to 3G, 4G, and 5G, respectively.

A foundational document in the definition of IMT-2020 is *ITU-R Recommendation M.2083*<sup>[1]</sup>. In broad strokes, this document develops a vision of the 5G mobile broadband connected society and future IMT. The two main contributions of this recommendation are a set of target values for key capabilities and a definition of usage scenarios, discussed subsequently.

M.2083 lists the eight key capabilities for IMT, together with the minimum requirements for each. The objectives that determined these target values follow:


- The user experience with IMT-2020 using a mobile device should match—to the extent possible—the experience with fixed networks. The enhancement will be realized by increased peak and user experienced data rate, enhanced spectrum efficiency, reduced latency, and enhanced mobility support.
- IMT-2020 should support massive machine-to-machine interconnections, for a variety of IoT environments.
- IMT-2020 should be able to provide these capabilities without undue burden in terms of energy consumption, network equipment cost, and deployment cost to make future IMT sustainable and affordable.

Figure 2, from M.2083, compares the capability requirements of IMT-2020 (5G) to those of IMT-Advanced (4G). It shows that substantial improvements are mandated for all eight capabilities, with the most substantial required improvements in the areas of traffic capacity and network energy efficiency.

The target values were published in 2015, with the admonition that they are presented for purposes of research and development and may be revised in light of future studies and implementation experience. This list was expanded and refined in 2017 into 13 technical performance requirements in ITU-R Report M.2410<sup>[2]</sup>. The purpose of these performance requirements is to assure that there should be a noticeable improvement of user *Quality of Experience* (QoE) for legacy 4G services and applications, and a high QoE for emerging 5G services and applications. Two terms should be distinguished:

- *Quality of Service* (QoS): The measurable end-to-end performance properties of a network service, which can be guaranteed in advance by a *Service-Level Agreement* (SLA) between a user and a service provider, so as to satisfy specific customer application requirements. Note: These properties may include throughput (bandwidth), transit delay (latency), error rates, priority, security, packet loss, and packet jitter.
- *Quality of Experience* (QoE): A subjective measure of performance in a system. QoE relies on human opinion and differs from QoS, which you can measure precisely.

Figure 2: Enhancement of Key Capabilities from IMT-Advanced to IMT-2020



In essence, the performance requirements for 5G are QoS measures designed to produce a high QoE. The M.2410 minimum technical performance requirements are as follows:

- *Peak Data Rate*: The maximum achievable data rate under ideal conditions per user/device (in Gbps). The minimum target values are downlink peak data rate of 20 Gbps and uplink peak data rate of 10 Gbps.
- *Peak Spectral Efficiency*: The maximum data rate under ideal conditions normalized by channel bandwidth (in bits/s/Hz). Another way of expressing this term is that it is the maximum data rate that can be transmitted over a given bandwidth. The relationship can be expressed as follows:  $R_p = W \times SE_p$ , where  $R_p$  is peak data rate,  $W$  is the available bandwidth, and  $SE_p$  is peak spectral efficiency. The minimum for peak spectral efficiencies is a downlink of 30 bps/Hz and uplink of 15 bps/Hz.
- *User-Experienced Data Rate*: The achievable data rate that is available ubiquitously across the coverage area to a mobile user/device (in Mbps or Gbps). This rate will depend the type of environment.

- *5th Percentile User Spectral Efficiency*: The 5% point of the cumulative distribution function of the normalized user throughput. The normalized user throughput is defined as the number of correctly received bits, that is, the number of bits contained in the *Service Data Units* (SDUs) delivered to Layer 3, over a certain period of time, divided by the channel bandwidth; it is measured in bits/s/Hz.
- *Average Spectral Efficiency*: The average data throughput per unit of spectrum resource and per cell (bits/s/Hz). The goal is a spectral efficiency of three times higher than IMT-Advanced.
- *Area Traffic Capacity*: The total traffic throughput served per geographic area (in Mbps/m<sup>2</sup>).
- *Latency*: Deals with transmission delays introduced by the network. Report M.2410 considers two types of latency:
  - *User-Plane Latency*: The contribution by the radio network to the time from when the source sends a packet to when the destination receives it (in ms).
  - *Control-Plane Latency*: Refers to the transition time from a most “battery efficient” state (for example, Idle state) to the start of continuous data transfer (for example, Active State). The minimum requirement is 20 ms.
- *Connection Density*: The total number of connected and/or accessible devices per unit area (per km<sup>2</sup>) that fulfills a specific QoS. The minimum requirement is 106/km<sup>2</sup>.
- *Energy Efficiency*: In general terms, the relation between useful output and energy consumption. In the context of M.2410, this parameter has two aspects:
  - *Network Energy Efficiency*: Refers to the quantity of information bits transmitted to/received from users, per unit of energy consumption of the RAN (in bits/Joule). The objective is efficient data transmission when the load on the network is substantial. The energy consumption for the RAN of IMT-2020 should not be greater than for IMT-Advanced, while delivering the enhanced capabilities. The network energy efficiency should therefore be improved by a factor at least as great as the envisaged traffic capacity increase of IMT-2020 relative to IMT-Advanced.
  - *Device Energy Efficiency*: Refers to a quantity of information bits per unit of energy consumption of the communication module (in bits/Joule). The objective is low-energy consumption when no data is being sent or received.
- *Reliability*: The probability of successful transmission of a Layer 2/3 packet within a required maximum time, which is the time it takes to deliver a small data packet from the radio protocol Layer 2/3 SDU ingress point to the radio protocol Layer 2/3 SDU egress point of the radio interface at a certain channel quality.



- *Mobility*: the maximum speed at which a defined QoS and seamless transfer between radio nodes that may belong to different layers and/or radio-access technologies (multi-layer/-RAT) can be achieved (in km/h). The following classes of mobility are defined:
  - *Stationary*: 0 km/hr
  - *Pedestrian*: 0 to 10 km/hr
  - *Vehicular*: 10 to 120 km/hr
  - *High-Speed Vehicular*: 120 to 5000 km/hr
- *Mobility Interruption Time*: The smallest time delay the system supports, during which the end-user device cannot exchange packets with any base stations during transmissions. The mobility interruption time includes the time required to execute any RAN procedure, radio resource control signaling protocol, or other message exchanges between the mobile station and the RAN. The minimum requirement is 0 ms.
- *Bandwidth*: The maximum aggregated system bandwidth. The minimum requirement is 100 MHz.

#### ITU-T and IMT-2020 “Softwarization”

As mentioned previously, the role of ITU-T in defining requirements and developing standards for IMT-2020 is complementary to that of ITU-R. ITU-T; it specifies requirements for overall non-radio aspects of the IMT-2020 network, especially with respect to network operations and support of service requirements. ITU-T Recommendations cover the core, RAN, and transport networks. ITU-T Y.3101<sup>[3]</sup> lists the following objectives with respect to IMT-2020:

1. Minimized dependency on access network technologies
2. Coping with traffic explosion in urban areas
3. Easy incorporation of future emerging services
4. Provision of a cost-efficient infrastructure
5. Expansion of the geographic reach of the network

The ITU-T approach to achieving these objectives depends in large part on the introduction of *network softwarization* in IMT-2020 network components. Y.3101 defines network softwarization as an overall approach for designing, implementing, deploying, managing, and maintaining network equipment and/or network components by software programming. It enables you to use modular network functions that you can deploy and scale on demand to accommodate various use cases easily and cost-efficiently.

Four aspects of network softwarization are important in 5G networks and are reflected in the ITU-T documents:

- *Software-Defined Networking* (SDN): An approach to designing, building, and operating large-scale networks based on programming the forwarding decisions in routers and switches with software from a central server. SDN differs from traditional networking, which requires configuring each device separately and relies on protocols that you cannot alter.



- *Network Functions Virtualization (NFV)*<sup>[18]</sup>: The virtualization of compute, storage, and network functions by implementing these functions in software and running them on virtual machines. NFV decouples network functions, such as routing, firewalling, intrusion detection, and network address translation from proprietary hardware platforms and implements these functions in software. It uses standard virtualization technologies that run on high-performance hardware to virtualize network functions. It is applicable to any data-plane processing or control-plane function in both wired and wireless network infrastructures.
- *Edge Computing*: A distributed *Information Technology (IT)* architecture in which client data is processed at the periphery of the network, as close to the originating source as possible.
- *Cloud-Edge Computing*: A form of edge computing that offers application developers and service providers cloud computing capabilities, as well as an IT service environment, at the edge of a network. The aim is to deliver compute, storage, and bandwidth much closer to data inputs and/or end users.

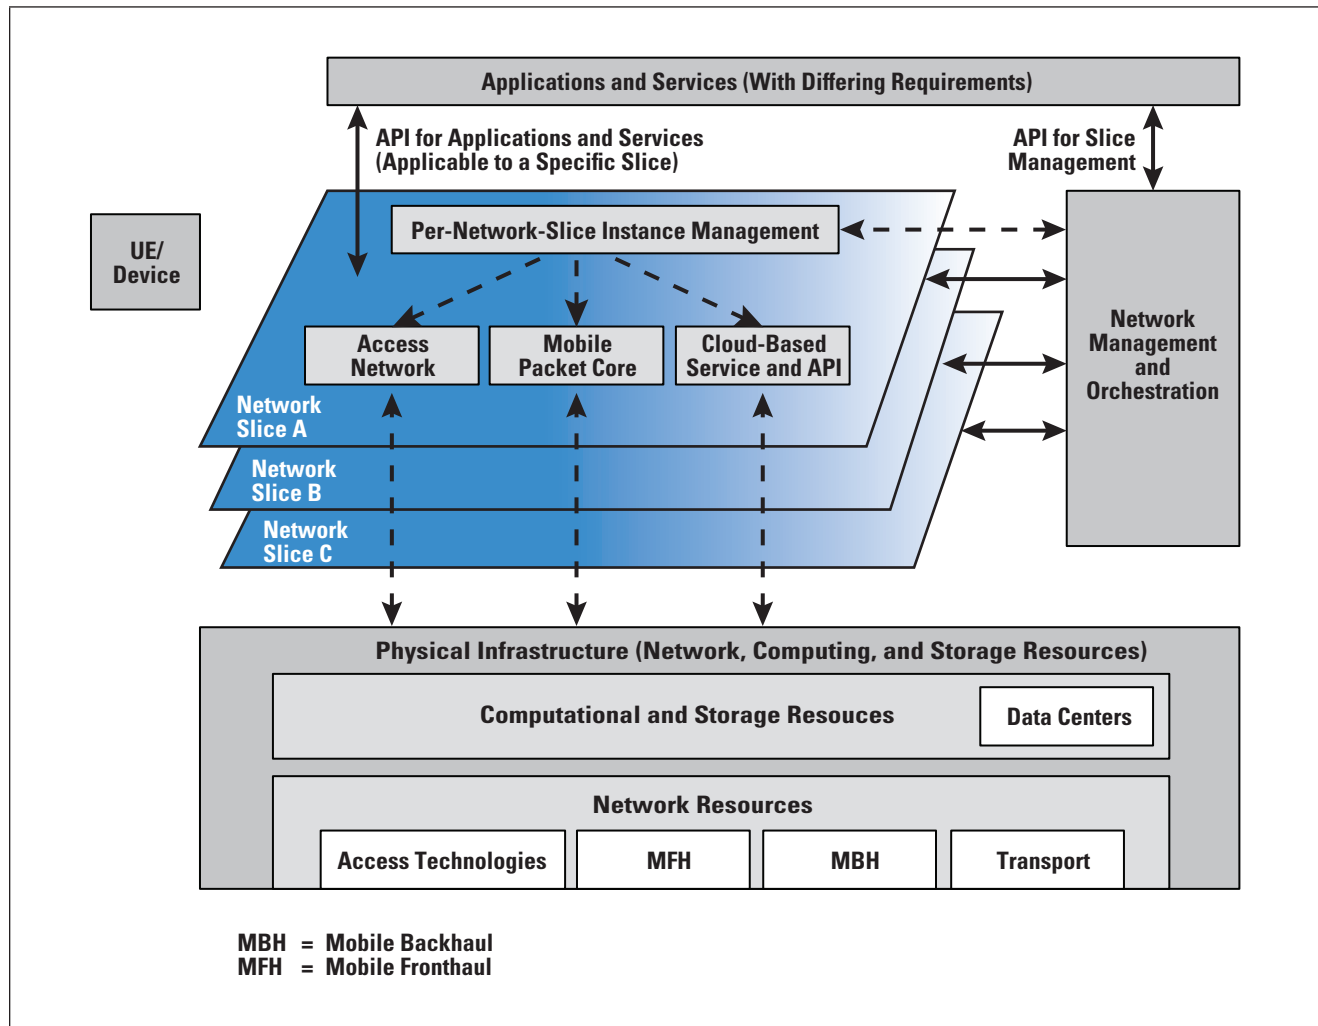
Network softwarization allows you to implement one of the essential features of 5G networks: *network slicing*. Network slicing permits you to separate a physical network into multiple virtual networks (logical segments) that can support different QoS requirements from applications and end users. Network slicing involves the selection and reservation of resources in the air interface, the RAN, the transport network, and the core network.

In essence, network slicing allows you to create multiple virtual networks atop a shared physical infrastructure. In this virtualized network scenario, physical components are secondary and logical (software-based) partitions are paramount, devoting capacity to certain purposes dynamically, according to your need. As your needs change, so can your devoted resources. Using common resources such as storage and processors, network slicing enables you to create slices devoted to logical, self-contained, and partitioned network functions. Network slicing supports the creation of virtual networks to provide a given QoS, such as guaranteed delay, throughput, reliability, and/or priority.

Figure 3, from ITU-T Recommendation Y.3150<sup>[4]</sup>, illustrates how network softwarization is incorporated in the design of IMT-2020 networks. The underlying physical infrastructure consists of a heterogeneous collection of network, computing, and storage resources. The figure shows four network resource categories. The access technologies consist of the resources at the air interface, including bandwidth, access protocol, channel coding, and modulation scheme. The mobile fronthaul refers to network paths between centralized radio controllers and remote radio units of a base-station function. The mobile backhaul refers to the network path between base-station systems and a core network. The transport resources consist of the switching hardware and software for routing data packets in the transport and core networks; an SDN controller manages these packets.

Using NFV, these underlying resources are abstracted to virtual resources used to create network slices, under the control of the management and orchestration function. Individual network slices can have specific characteristics that reflect various different requirements derived from application and services.

Figure 3: Network Softwarization for IMT-2020




**ITU-T IMT-2020 Core Network Framework**

ITU-T Recommendation Y.3102<sup>[5]</sup> provides the framework for overall non-radio aspects of the IMT-2020 network. Figure 4, from Y.3102, illustrates the interactions between the network functions for providing network service.

The framework delineates three domains. The *User Equipment* (UE) domain consists of devices that transmit and receive data over the IMT-2020 network. The *Access Network* (AN) domain is the wireless connection between the UE and the *Core Network* (CN), defined by the ITU-R radio interface recommendations.

The framework diagram also depicts the division between a control plane and a user plane, which cuts across the AN and CN.

Figure 4: Framework of the IMT-2020 Network



The *Control Plane* performs the call and connection control functions. For this purpose, a signaling connection between the UE and the CN exchanges signaling messages that manage the signaling connection and the call established for the UE. The Control-Plane functions are requested and managed via control signals that are exchanged between UE and the AN, and between the AN and the CN. Through signaling, the control plane sets up and releases connections, and may restore a connection if a failure occurs. The control plane also performs other functions that support call and connection control, such as routing information dissemination.

The core network includes the following functional elements:

- *Network Access Control Function (NACF)*: Provides access to the CN services for the AN and UE. NACF includes:
  - *Registration Management*: Enables UE to register for network access. NACF performs, but is not limited to, network slice instance selection, UE authentication, authorization of network access and network services, and network access policy control.
  - *Connection Management*: Establishes and releases a signaling connection between the UE and the core network.

- *Session Management Function Selection*: Determines the session management function that is most appropriate to establish and manage a session. In the context of IMT-2020, a session is an association between UE and a data network that provides a *Protocol Data Unit* (PDU) connectivity service.
- *Session Management Function* (SMF): Sets up and manages one or more sessions that provide connectivity between the local UE and a remote UE. This function deals with user path selection and enforcement of policies, including QoS policy and charging policy.
- *Policy Control Function* (PCF): Provides for control and management of policy rules.
- *Capability Exposure Function* (CEF): Enables the exposure of network functions and network slices as a service to third parties.
- *Network Function Registry Function* (NFR): Assists the discovery and selection of required network functions.
- *Unified Subscription Management Function* (USM): Stores and manages UE context and subscription information including, but not limited to, UE information on registration and mobility management, information on network functions that serve the UE, and information on session management. USM also provides UE authentication information to the *Authentication Server Function* (ASF).
- *Network Slice Selection Function* (NSSF): When UE requests registration with the network, NSSF sends a network slice selection request to NSSF with preferred network slice selection information. The NSSF responds with a message including the list of appropriate network slice instances for the UE.
- *Authentication Server Function* (ASF): Performs authentication between UE and the network.
- *Application Function* (AF): Interacts with application services that require dynamic policy control. AF extracts session-related information (for example, QoS requirements) from application signaling and provides it to PCF in support of its rule generation.

The user plane refers to the set of traffic forwarding components through which traffic flows. Its principal function is to provide transfer of end-user information.

The sole functional element in the user plane is the *User Plane Function* (UPF). This function includes traffic routing and forwarding, *Protocol Data Unit* (PDU) session tunnel management, and QoS enforcement. The PDU session tunnels are used between AN and UPF(s) as well as between different UPFs as user-plane data transport for PDU sessions. UPF also provides optional functions including packet inspection and collection of *User-Plane* (UP) traffic for lawful intercept. In order to accommodate the diversity of network scenarios, UPF may also provide interworking functions among different network segments, for example, interworking between the IP-based core network and the non-IP-based access network.

Y.3102 also lists the primary network services that the supported core network framework supports. They include:

- *Registration Management (RM)*: Register or deregister UE with the IMT-2020 network and establish the user context in the network.
- *Connection Management (CM)*: Establish and release the signaling connection between the UE and NACF.
- *Session Management (SM)*: Manages PDU sessions including control of PDU session tunnel establishment, modification, and release.
- *User-Plane Management (UPM)*: Forward user traffic, including user traffic rerouting between UPFs because of the serving UPF relocation and enforcement of QoS policies.
- *Mobility Management (MM)*: Used to handle all aspects related to UE mobility. Mobility management aspects include, but are not limited to, UE reachability and handover management.

### 3GPP

The *3rd Generation Partnership Project (3GPP)* was formed in 1998 by a global consortium of regional *Standards Development Organizations (SDOs)* to develop technology specifications for 3G cellular networks. Because it involved the efforts of the world's leading national standards organizations, 3GPP became the dominant agent in the development of specifications for 3G, then 4G, and now 5G cellular networks.

3GPP began work in 2016 on defining 5G technical specifications for a new radio access technology, known as 5G NR (New Radio) and a next-generation network architecture (5G NextGen). Unlike previous generations, competing standards bodies are no longer working on potential solutions for 5G.

Figure 5 shows the key players in the 3GPP process and their relationships to one another. Within the 3GPP organization is a *Project Coordination Group (PCG)*. It is responsible for overall time frame and management of technical work to ensure that the 3GPP specifications are produced in a timely manner as required by the marketplace. Subordinate to the PCG are three *Technical Specification Groups (TSGs)*. Each TSG has the responsibility to prepare, approve, and maintain the specifications within its terms of reference; it may organize its work in *Working Groups (WGs)* and liaise with other groups as appropriate. The TSGs report to the PCG.

Key to the 3GPP process are the *organizational partners*. An organizational partner is a standards organization with a national, regional, or other officially recognized status (in its country or region) that has the capability and authority to publish standards nationally or regionally. Associated with organizational partners are individual members, which are member companies affiliated with one of the organizational partners.

Finally, there are *market-representation partners*, which are organizations invited to participate by the organizational partners to offer market advice to 3GPP and to bring into 3GPP a consensus view of market requirements (for example, services, features, and functions) falling within the 3GPP scope.

Figure 5: 3GPP Process

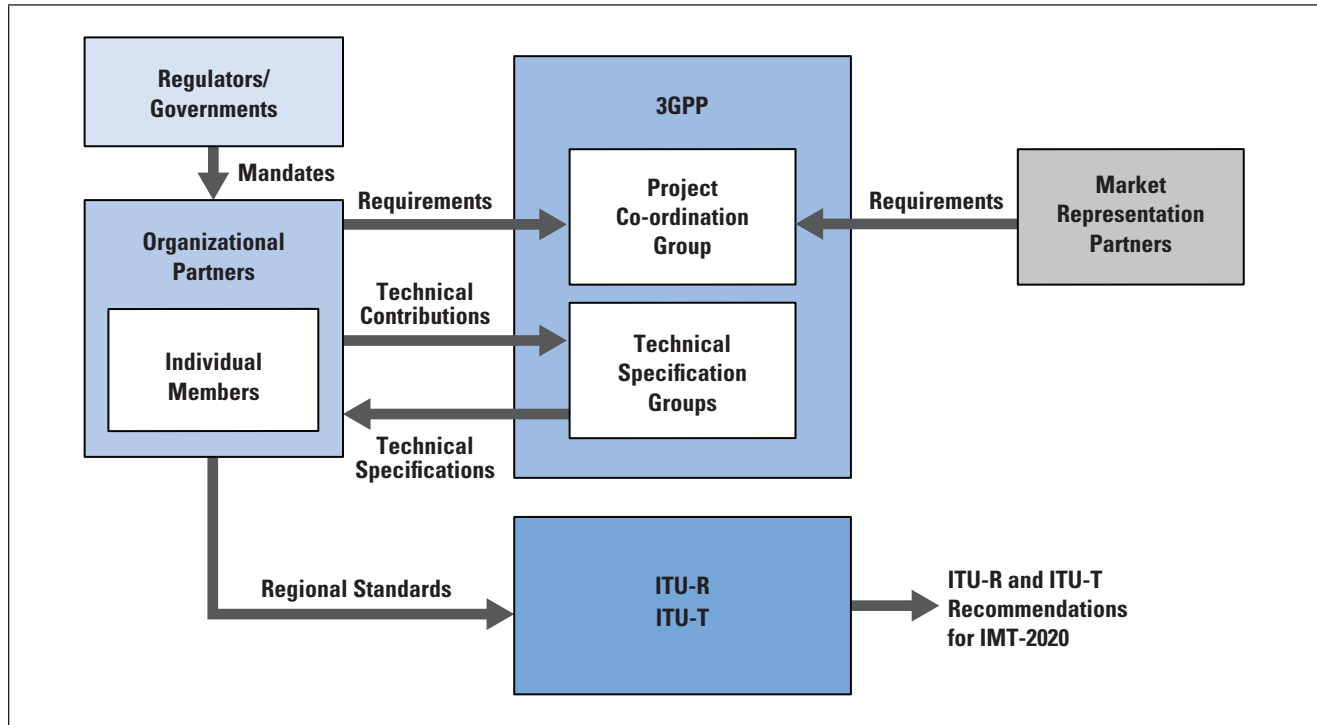


Figure 6 shows, in general terms, the flow of information between the previously mentioned entities. The PCG plans the work of 3GPP based on requirements provided by the organizational partners and the market representation partners. The organizational partners are influenced particularly by their respective national and regional governments and regulators, whereas the market representation partners generate requirements dictated by the potential market. Individual members provide technical contributions to the TSGs, which ultimately result in technical specifications. These specifications are transmitted from the TSGs to the organizational partners, who translate them into national and regional standards. Finally, these standards serve as input to ITU in the development of 5G-related Recommendations.

### 3GPP Releases

3GPP uses a system of parallel *Releases* that provide developers with a stable platform for the implementation of features at a given point and then allow for the addition of new functions in subsequent Releases. Releases are staggered and work is done on multiple Releases in parallel at different stages. When a Release is finalized, it means that all new features are functionally frozen and ready for implementation. Furthermore, each 3GPP Release is self-contained, meaning that you can build a cellular system based on the set of frozen specifications in that Release.

As such, Releases do not just contain the newly implemented features, but instead are introduced in a highly iterative manner that builds upon previous Releases. Table 1 provides information on the three releases relating to 5G that are completed at the time of this writing. Release 15 provided an early definition of useful 5G features to enable deployment by 2020. Subsequent releases add progressively more functions. Release 16 should closely resemble the initial set of IMT-2020 Recommendations issued by ITU in 2020.

Table 1. 3GPP Releases for 5G

Release #	Status	Functional Freeze	End Date
Release 17	Frozen	2022-03-18	2022-06-10
Release 16	Frozen	2020-07-03	2020-07-03
Release 15	Frozen	2019-03-22	2019-06-07

When a Release is frozen, the TSGs can add no additional functions to the specifications. However, detailed protocol specifications may not yet be complete. The end date shown in Table 1 is indicative only, since for each Release, a considerable number of refinements and corrections can be expected for at least two years following this date.

### 3GPP Requirements for 5G

The 3GPP documents include a description of 5G requirements that are significantly more detailed than those provided in the ITU documents. As such, they provide an important guide to implementers of 5G networks, components, and systems as to what the market requirements are for 5G success (Refer to Figure 6).

Figure 6: 3GPP Basic Capability Requirements

Network Slicing Diverse Mobility Magement Multiple Access Technologies Resource Efficiency Efficient User Plane Efficient Content Delivery Priority, QoS, and Policy Control Dynamic Policy Control Connectivity Models Network Capability Exposure Context Aware Network Self-Backhaul Flexible Broadcast/Multicast Service	Subscription Aspects Energy Efficiency Markets Requiring Minimal Service Levels Extreme Long-Range Coverage in Low-Density Areas Multi-Network Connectivity and Service Delivery Across Operators 3GPP Access Network Selection eV2X Aspects NG-RAN Sharing Unified Access Control QoS Monitoring Ethernet Transport Services	Non-Public Networks 5G LAN-Type Service Positioning Services Cyber-Physical Control Applications in Vertical Domains Messaging Aspects Steering of Roaming Minimization of Service Interruption UAV Aspects Video, Imaging, and Audio for Professional Applications Critical Medical Applications
--	---	--

eV2X = Enhanced Vehicle-to-Everything

UAV = Unmanned Aerial Vehicle

3GPP Technical Specification TS 22.261<sup>[6]</sup> defines requirements for 34 basic capabilities to be provided by a 5G network; they are listed in Figure 6. For each capability, TS 22.261 provides a description and elaborates on the requirements for that capability.



TS 22.261 also lists performance requirements that are more detailed and more demanding than those defined in ITU-R Report M.2410. The requirements cover the following categories:

- *High Data Rates and Traffic Densities*: Several 5G scenarios require the support of very high data rates or traffic densities, including urban and rural areas, office and home, and special deployments (for example, massive gatherings, broadcast, residential, and high-speed vehicles).
- *Low Latency and High Reliability*: Some scenarios require the support of very low latency and very high communications service availability, which in turn implies very high reliability. The overall service latency depends on the delay on the radio interface, transmission within the 5G system, transmission to a server that may be outside the 5G system, and data processing. Some of these factors depend directly on the 5G system itself, whereas for others the impact can be reduced by suitable interconnections between the 5G system and services or servers outside of the 5G system, for example, to allow local hosting of the services. TS 22.261 provides an overview of potential scenarios and references other technical specifications for specific requirements.
- *High Accuracy Positioning*: The 5G System shall provide different 5G positioning services with configurable performances working points (for example, accuracy, positioning service availability, positioning service latency, energy consumption, update rate, and time to first fix) according to the needs of users, operators, and third parties. TS 22.261 lists quantitative requirements for numerous indoor and outdoor scenarios.
- *Key Performance Indicators (KPIs) for a 5G System with Satellite Access*: In some contexts, a 5G access network will use at least one satellite link. KPIs defined in TS 22.261 include minimum and maximum UE-to-satellite delay for various earth orbits, as well as maximum propagation delay.
- *High Availability IoT Traffic*: This requirement is concerned specifically with medical monitoring but is applicable to other scenarios that require highly reliable machine-type communication in both stationary and highly mobile settings.
- *High Data Rate and Low Latency*: This requirement defines data and latency requirements for such scenarios as audio-visual interaction, gaming, and virtual reality.
- *KPIs for UE-to-Network Relaying in 5G System*: In several scenarios, it can be beneficial to relay communication between one UE and the network via one or more other UEs. This category includes performance requirements for various scenarios.


#### Usage Scenarios and Use Cases

Two important concepts in ITU-R Recommendation M.2083 and related documents are *Usage Scenario* and *Use Case*. No ITU document defines these terms, but the following definitions should suffice for this article.

- *Usage Scenario*: A general description of a way in which an IMT network is used. A usage scenario dictates various performance and technical requirements. A wide but nevertheless constrained variety of use cases are encompassed by a usage scenario.
- *Use Case*: A specific application or way of using an IMT network; also, a general account of a situation or course of actions that use an IMT network. It is described from the end user's perspective and illustrates fundamental characteristics. A use case dictates more specific and refined performance and technical requirements than the corresponding usage scenario.

M.2083 defines three usage scenarios: enhanced mobile broadband, massive machine-type communications, and ultra-reliable and low-latency communications. Figure 7, from M.2083, indicates the relative importance of the key capabilities for the three usage scenarios.

Figure 7: The Importance of Key Capabilities in Different Usage Scenarios



### Enhanced Mobile Broadband

*Enhanced Mobile Broadband (eMBB)* is the 5G feature that provides a significant increase in data rate over 4G for a normal mobile Internet user. Enhanced mobile broadband services allow users to experience high-speed and high-quality multimedia services such as virtual reality, *Augmented Reality (AR)*, and 4,000-pixels horizontal resolution (4K) video, at any time and place.

These applications require reasonably low-latency and good connection density, with a high demand on the other six key capabilities. In addition to the consumption of multimedia content for entertainment purposes, eMBB supports numerous business applications. They include cloud access apps for commuters and other off-site employees, the ability of remote workers to communicate with the back office, or indeed an entire smart office where all devices are wirelessly and seamlessly connected.

Of the three usage scenarios defined in M.2083, eMBB is the only general-purpose case, and it is the one that is most familiar to current 4G users. In essence, eMBB is an enhanced version of 4G, providing improved performance and an increasingly seamless user experience.

ITU-R Report M.2410 lists three deployment options that characterize the scope of eMBB and that are used for purposes of evaluation of candidate specifications: indoor hotspot, dense urban, and rural. The remainder of this section provides a brief overview of all three.

ITU-R Report M.2412<sup>[7]</sup> defines *Indoor Hotspot* as “...an indoor isolated environment at offices and/or in shopping malls based on stationary and pedestrian users with very high user density.” This deployment scenario focuses on small coverage per site/*Transmission and Reception Point* (TRxP) and high user throughput or user density in buildings. The key characteristics of this deployment scenario are high capacity, high user density, and consistent user experience indoors.

5G capabilities should enable a seamless interface for users moving into and out of the indoor zone, without the necessity of joining a Wi-Fi network for indoor use. Types of demand include frequent upload and download of data from a company’s servers and real-time video meetings with local as well as remote participants.

One of the main challenges for supporting 5G use cases in the indoor environment is a consequence of the use of much higher-frequency bands for 5G than are used for 4G and earlier generations. These higher bands lead to greater link losses. For example, outdoor signals on the C band will be subject to an 8- to 13-dB link loss when penetrating through one concrete wall. The signals on the higher-mm wave band will experience difficulty in penetrating through a wall as the link loss exceeds 60 dB. It is a considerable challenge for outdoor 5G macro signals to cover indoor areas, and a dedicated 5G network consisting of interconnected base stations will be required for indoor environments.

An example use case in the indoor hotspot category is the smart office. The installation of 5G networks in the office environment can enable dramatic changes in the capabilities that businesses can exploit.

Examples of features now in use or that may soon be in use in 5G-enabled workplaces include:

- Facial recognition can be used for entrance security. The employee need not carry an identification tag or use some sort of token to gain entrance.
- A 5G virtual desktop infrastructure enables workers to connect their mobile device on a docking pad to a cloud computing system.
- Workers can convene remote conferences, talking to each other's avatars in cyberspace.
- Security systems can use high-definition video to monitor in greater detail and expand the ability to scan for security threats.
- Workers have faster access to a broader selection of apps.
- 5G enables real-time collaboration between people and things, possibly including augmented reality features.
- Real-time video interaction will become standard. This access allows capabilities such as real-time troubleshooting and ad hoc meetings.
- Synchronization of local data with the cloud becomes almost instantaneous, further enhancing collaboration.
- Sensors or facial recognition can tell if people are in the building and where they are at any given moment.

In essence, the smart office use case is characterized by heavy data use, with a particular reliance on high-definition video, in an indoor environment with low mobility requirements. In this use case scenario hundreds of users require ultra-high bandwidth to serve intense bandwidth applications. To some extent, Wi-Fi supports these capabilities, but with the increasing demands for high traffic volume, high density of users, and seamless integration of local and wide-area communications, a unified 5G solution has inherent advantages over a mixed Wi-Fi/cellular environment.

ITU-R Report M.2412<sup>[7]</sup> defines *Dense Urban* as "...an urban environment with high user density and traffic loads focusing on pedestrian and vehicular users." The dense urban microcellular deployment scenario focuses on macro TRxPs with or without micro TRxPs and high user densities and traffic loads in city centers and dense urban areas. The key characteristics of this deployment scenario are high traffic loads and outdoor-to-outdoor and outdoor-to-indoor coverage.

The dense urban environment for 5G is characterized by the use of a dense collection of small cells to supplement macro cells for two reasons<sup>[8]</sup>:

- The concentrated collection of stationary, pedestrian, and vehicular users, with 5G use cases, generates a tremendous traffic load.
- 5G-mm Wave networks are predominantly noise-limited. The result is that only small cell sizes can be supported.

An example use case in this category is provided by the EU project METIS (*Mobile and wireless communications Enablers for the Twenty-twenty Information Society*)<sup>[9]</sup>. It refers to the connectivity and data rates required for users of high-volume services at any place and at any time in a dense urban environment, including both user interaction with cloud services and data- and device-centric services.

5G enables enhanced cloud services beyond the traditional services of web browsing, file download, and social media. Enhanced services include high-definition video streaming and video sharing. Enhanced device-centric services include augmented reality with information fetched from sensors, smart phones, wirelessly connected cameras, and other sources. The main features of this use case follow:

- High traffic loads
- Low mobility
- High data rate
- Outdoor coverage
- Outdoor-to-indoor coverage
- Support for both low and high frequency
- Limited interference
- High user density

This use case presents two unique challenges:

- Users expect the same QoE in any context, including at their workplace, enjoying leisure activities such as shopping or being on the move, on foot, or in a vehicle.
- Users in urban environments tend to dynamically cluster. Examples include people waiting at a traffic light or bus stop and conference room meetings at the workplace. These clusters lead to sudden peaks of geographically concentrated mobile broadband demand.

ITU-R Report M.2412<sup>[7]</sup> defines *Rural-eMBB* as “...a rural environment with larger and continuous wide area coverage, supporting pedestrian, vehicular and high-speed vehicular users.” The rural deployment scenario focuses on larger and continuous coverage. The key characteristics of this scenario are continuous wide area coverage supporting high-speed vehicles. This scenario uses macro TRxPs, and is noise- and/or interference-limited.

The rural deployment also supports last-mile service to residences and other subscribers to provide telephone and Internet access. Many homes may be near a fiber connection, but the deployment of the last mile of the cabling can be very expensive and not necessarily cost-effective. The addition of new subscribers, or households, may be very expensive if new cables need to be installed. It may also require the operator to support two distinct systems, each with its own subscription management, for wired and wireless subscribers. To address this problem, delivering the last mile wirelessly may be a viable option. Such solutions are known as *Wireless Local Loop* (WLL), where the last mile is delivered wirelessly.

### Massive Machine Type Communications

*Massive Machine Type Communications* (mMTC) is characterized by a very large number of connected devices typically transmitting a relatively low volume of non-delay-sensitive data. However, the machine-to-machine communications involves a range of performance and operational requirements. Devices are required to be low-cost and have a very long battery life, such as five years or longer.

The mMTC usage scenario defined by ITU-R represents a subset of the total IoT universe. A white paper from Ericsson<sup>[10]</sup> lists four segments that comprise IoT:

- *Massive IoT*: Massive IoT is characterized by huge volumes of constrained devices that send and/or receive messages infrequently. The traffic is often tolerant of delay. Examples of use cases include low-cost sensors, meters, wearables, and trackers. Such devices are often deployed in challenging radio conditions such as in the basement of a building. Therefore, they require extended coverage and may rely solely on a battery power supply that puts extreme requirements on the life of the battery.
- *Broadband IoT*: Broadband IoT is an application of eMBB to the IoT environment, providing high data rates and relatively low latencies. Examples of use cases are in the areas of automotive, drones, *Augmented Reality/Virtual Reality* (AR/VR), utilities, manufacturing, and wearables.
- *Critical IoT*: Critical IoT is an application of *Ultra Reliable and Low Latency Communications* (URLCC) to the IoT environment, providing extremely low latencies and ultra-high reliability at a variety of data rates. In contrast to Broadband IoT, which achieves low latency on best effort, critical IoT is intended to deliver data within strict latency bounds with required guarantee levels, even in heavily loaded networks. Examples of use cases are in the areas of intelligent transportation systems, smart utilities, remote healthcare, smart manufacturing, and fully immersive AR/VR.
- *Industrial Automation IoT*: This segment supports seamless integration of cellular connectivity into the wired industrial infrastructure used for real-time advanced automation. These applications have extremely demanding requirements such as very accurate indoor positioning and time synchronization across devices and networks.

Massive IoT, as defined by Ericsson, is equivalent to mMTC defined by ITU-R. In terms of the number of connections, mMTC is the most rapidly growing segment of IoT<sup>[11]</sup>. Table 2, based on a 2020 Ericsson white paper<sup>[12]</sup>, indicates likely mMTC use cases that 5G supports.

Table 2: Industry and Society Applications Enabled by Massive IoT

Application Area	Use Cases
Transport and Logistics	Fleet Management Goods Tracking
Agriculture	Climate / Agriculture Monitoring Livestock Tracking
Environment	Process Monitoring and Control Maintenance Monitoring
Industrial	Process Monitoring and Control Maintenance Monitoring
Utilities	Smart Metering Smart Grid Management
Smart Cities	Parking Sensors Smart Bicycles Waste Management Smart Lighting
Smart Buildings	Smoke Detectors Alarm Systems Home Automation
Consumers	Wearables Children/Elderly Tracking Medical Monitoring

An important group of mMTC use cases are in the general category of smart cities. ITU-T 4900<sup>[13]</sup> defines a *Smart Sustainable City*, or simply *Smart City*, as follows: A smart sustainable city is an innovative city that uses *Information and Communication Technologies* (ICTs) and other means to improve quality of life, efficiency of urban operation and services, and competitiveness, while ensuring that it meets the needs of present and future generations with respect to economic, social, environmental, and cultural aspects.

The sustainability of a smart city is based on four main aspects:

- *Economic*: The ability to generate income and employment for the livelihood of the inhabitants.
- *Social*: The ability to ensure that the welfare (safety, health, education) of the citizens can be equally delivered despite differences in class, race, or gender.
- *Environmental*: The ability to protect future quality and reproducibility of natural resources.
- *Governance*: The ability to maintain social conditions of stability, democracy, participation, and justice.

Some smart-city use cases, such as *Public Protection and Disaster Relief* (PPDR), fit into the URLCC usage scenario, but many others fall into the mMTC usage scenario category.



As examples, *ITU-T Series Y Supplement 56*<sup>[14]</sup> lists the following mMTC demonstration examples of smart-city use cases:

- *Pedestrian Monitoring for Decisive Disaster Response*: Involves the installation of surveillance cameras throughout a city that can monitor crowd size and behavior and transmit this information to a central monitoring/management source. The cameras monitor locations that are likely to draw large groups of people, such as near a railroad or subway station or near a school. If a disaster occurs near one of these sites, the system provides real-time information about the size of the crowd at risk. In addition, the pedestrian monitoring system facilitates the understanding of the behavior of crowds and the detection of abnormal situations by analyzing images captured by surveillance cameras. If it detects any abnormality, the system automatically provides information or instructions for evacuation from the disaster site or for prevention of accidents.
- *Citizens' Safety Services*: Involves interworking between smart-city operation centers and fire and police stations for the citizens' safety services. Surveillance cameras are deployed throughout the city to provide extensive coverage with a minimum number of cameras. IoT-enabled traffic sensors deployed throughout the city can measure rate and volume of traffic. The operations center connects wirelessly to the cameras and sensors to provide a central source of information. The operations center provides traffic information to the first responders to enable them to take the best route to the scene of an emergency.
- *Lift Monitoring Services*: Involves monitoring lifts, or elevators, throughout a city. One such system developed by Surbana Jurong is deployed in Singapore and other Asian cities. The system consists of a central *Lift Monitoring System* (LMS) and IoT-enabled sensors installed in lifts throughout the city. The installation in Singapore monitors more than 26,000 lifts across 10,000 housing units. The system enables rapid response to elevator malfunction. In addition, the sensor devices capture data on an ongoing basis. This data, using machine-learning algorithms, is used to predict future failures, allowing for optimized maintenance and reduced downtime.
- *Infrastructure Monitoring*: Involves using IoT sensor devices to monitor aging infrastructure elements to support automated inspection, diagnosis, confirmation of repair effort, and subsequent status check. The scheme can be applied to bridges, tunnels, and paved roads.
- *Citizen Identification System Using Biometric*: Has objective to provide a digital identity to the entire population to serve as the basis for accessing social services and interacting with the government at various levels. One such system, called *Aadhaar*, is deployed nationwide in India and currently has over one billion people registered. In any large Indian city, there are tens or even hundreds of thousands of Aadhaar devices in use for registration and service access. These devices form a massive IoT network connected to a central server.

**Ultra-Reliable and Low-Latency Communications**

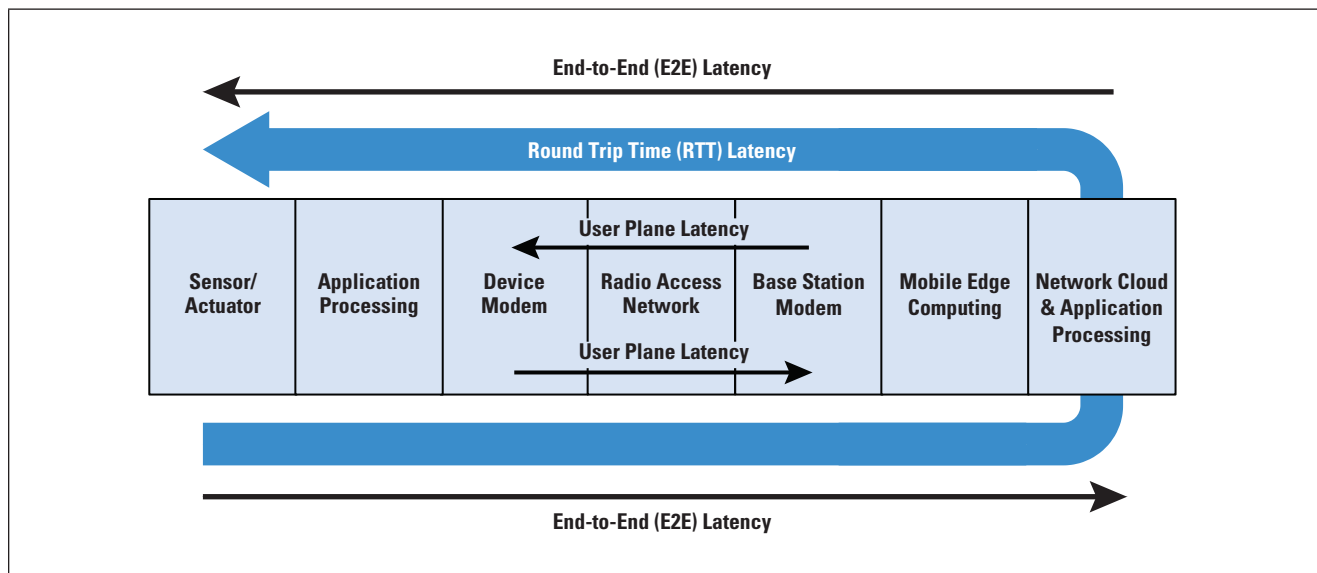
URLLC is a form of machine-to-machine communications that enables delay-sensitive and mission-critical services that require very low end-to-end delay, such as tactile Internet, remote control of medical or industrial robots, driverless cars, and real-time traffic control.

Figure 7 (earlier in this article) indicates that two parameters are of high importance for URLLC: *latency* and *mobility*. ITU-R Report M.2410 breaks the latency requirement into two parts:

- *User-Plane Latency*: is the contribution by the radio network to the time from when the source sends a packet to when the destination receives it (in ms). It is defined as the one-way time it takes to successfully deliver an application layer packet/message from the radio protocol Layer 2/3 SDU ingress point to the radio protocol Layer 2/3 SDU egress point of the radio interface in either uplink or downlink in the network for a given service in unloaded conditions, assuming the mobile station is in the active state. The minimum requirement (that is, the maximum allowable value) is 1 ms assuming unloaded conditions (that is, a single user) for small IP packets (for example, 0-byte payload + IP header), for both downlink and uplink.
- *Control-Plane Latency*: refers to the transition time from a most battery efficient state (for example, Idle state) to the start of continuous data transfer (for example, Active state). The minimum requirement is 20 ms.

User-Plane latency, however, is only one component that UE experiences overall, as illustrated in Figure 8. The *End-to-End (E2E)* latency is generally defined as the time it takes from when a data packet is sent from the transmitting end to when it is received at the receiving entity; for example, Internet server or other device. The measurement reference is the interface between Layers 2 and 3. It is also referred to as *One-Trip Time (OTT)*. It includes the user-plane latency in one direction, transport network delays, and application processing time.

Figure 8: E2E Latency and Round-Trip Time Latency



A related measure is *Round-Trip Time* (RTT), which is the time from when a data packet is sent from a source device until an acknowledgement or response is received from the destination device. Unfortunately, E2E latency is sometimes equal to RTT latency in the literature, even in some 3GPP documents. However, the implication in most standards and specification documents is that E2E latency refers to one-way latency, not round-trip.

Mobility is the maximum UE speed (in km/h) at which a QoS can be achieved. Mobility assumes a seamless transfer between radio nodes that may belong to different layers and/or radio access technologies (multi-layer/RAT) can be achieved. The following classes of mobility are defined:

- *Stationary*: 0 km/h
- *Pedestrian*: 0 to 10 km/h
- *Vehicular*: 10 to 120 km/h
- *High-speed vehicular*: 120 to 500 km/h

M.2410 does not provide a specific measure of QoS. Report ITU-R M.2412 defines QoS as successful delivery of 99% of messages within 10s.

Another aspect of mobility addressed in M.2410 is *Mobility Interruption Time*, which is the shortest time duration supported by the system during which UE cannot exchange user-plane packets with any base station during transitions. This number includes the time required to execute any RAN procedure, radio resource control signaling protocol, or other message exchanges between the mobile station and the RAN. The minimum requirement for mobility interruption time is 0 ms. Thus, there should be no interruption of service when moving UE switches from one base station to another.

#### URLLC Use Cases

A URLLC white paper from 5G Americas<sup>[15]</sup>, one of the 3GPP market representation partners, provides a useful way of understanding the wide variety of URLLC use cases by focusing on emerging mission-critical applications that have demanding reliability and latency requirements. These use cases include:

- Smart Factory
- Ground Vehicles, Drones, and Robots
- Tactile Interaction
- Augmented Reality and Virtual Reality
- Emergency, Disasters, and Public Safety
- Urgent Health Care
- Intelligent Transportation

The area that has perhaps received the most attention as an application area that requires URLLC support is that of the *Smart Factory* or *Industrial Automation*. This application area is typified by extremely demanding reliability and latency requirements for 5G communication links between sensors, actuators, and controllers.

Traditionally, Ethernet has been used to provide network connectivity. For smart factories, wireless networks provide many advantages over Ethernet:

- Reduced cost of manufacturing, installation, and maintenance
- Higher long-term reliability as wired connections suffer from wear and tear in motion applications
- Inherent deployment flexibility

With 5G, dispersed IoT sensors, actuators, controllers, and robots driven by software command and control can expand the ability to more fully automate an industrial process.

The application area that encompasses *Ground Vehicles, Drones, and Robots* refers to remotely controlled mobile devices and robots. Such devices are in common use in factory applications, but are also deployed in other contexts, such as smart agriculture. One area of particular interest is unmanned aircraft traffic management.

*Tactile Interaction* refers to a level of responsiveness that works at a human scale. For example, remote health care or gaming applications may require very low round-trip times to convince human senses that the perceived touch, sight, and sound are lifelike. These use cases involve interaction between humans and systems, where humans wirelessly control real and virtual objects, and the interaction requires a tactile control signal with audio or visual feedback. Robotic controls and interaction include several scenarios with many applications in manufacturing, remote medical care, and autonomous cars. The tactile interaction requires real-time reactions on the order of a few milliseconds. Remote surgery, discussed later in this article, is perhaps the most demanding use case. Table 3 gives typical values of *Key Performance Indicators* (KPIs) for tactile Internet applications.

Table 3: Key Performance Indicators for Tactile Internet

KPI	Value
Traffic Volume Density	0.03–1 Mbps/m <sup>2</sup> / (cell radius 100 m <sup>2</sup> )
Experienced User Throughput	0.3–1 Mbps (UL)
Latency	User-plane latency less than 2 ms
Availability	>99.999%
Reliability	>99.999 % for healthcare or remote driving/manipulation 95 % for remote gaming or remote augmented reality

AR and VR tend to have relatively high data-rate requirements. Some specific use cases also have URLLC requirements.

A paper from the *Next Generation Mobile Networks Alliance* (NGMN)<sup>[16]</sup> lists three AR/VR examples with URLLC requirements:

- *Augmented Worker*: Augmented work is work that integrates digital technologies into the industrial environment to improve how work is done. Augmented work is appropriate for situations when it is not cost-effective or even possible to fully automate tasks, but it is desirable to augment the capabilities of the human worker. A good example is a task such as equipment repair where the access is difficult (for example, a hazardous environment) or the expert is at a remote place. The remote worker can be equipped with an AR headset and some sort of tactile interface for remote control. Sensor information from the remote target location in terms of audio, video, and haptic (tactile) enables the remote operator to control actuators at the target location to achieve the required work.
- *360 Panoramic VR View Video Broadcasting*: 360-degree videos are video recordings where a view in every direction is recorded at the same time, shot using an omnidirectional camera or a collection of cameras. With 360 panoramic VR view video broadcasting, the video is broadcast in real time. Remote users with VR headsets can view the live video feed, and by turning their head, see the point of view change in real time.
- *AR and MR Cloud Gaming*: A good example of an application in the AR/VR area that requires URLLC performance is AR and *Mixed Reality* (MR) cloud gaming, which is real-time game playing using a thin client with the bulk of the software on edge servers. This online gaming service provides on-demand streaming of games onto computers, consoles, and mobile devices. Thus, the user does not have to upgrade frequently and to deal with compatibility issues. Highly interactive games with tight QoS requirements generate the need for low-latency network performance.

Use cases in the category of *Emergency, Disasters, and Public Safety* generally require high reliability to enable response to natural disasters and emergencies. Accurate position location and very low latency to enable rapid response are also often critical requirements.

The *Urgent Health Care* category refers to applications involving remote diagnosis and treatment. A white paper from 5G Americas<sup>[17]</sup> lists the following examples in this category:

- *Remote Patient Monitoring*: This use case involves remote patient monitoring via communication with devices that measure certain health indicators, such as pulse, blood glucose, blood pressure, and temperature. On an individual basis, the data rate and latency requirements are modest. However, for this use case to become pervasive, 5G is needed to support the massive increase in the number of connections per square meter while still maintaining the requisite QoS.

- *Remote Health Care*: This use case provides for individualized consultation, treatment, and patient monitoring built on a video linkup capability. The video conferencing can be augmented with remote transfer of health-related data in real time. Treatment could also be offered using smart pharmaceutical devices that correctly administer approved dosages of a drug on a schedule specified by the physician or practitioner.
- *Remote Surgery*: More demanding is remote surgery via control of robotic devices. This application area may be appropriate in ambulances, disaster sites, and remote areas. Important requirements are precise control and very low latency, very high reliability, and tight security.

The *Next Generation Mobile Networks Alliance*<sup>[16]</sup> lists the following examples in the Intelligent Transportation category:

- *Assisted Driving*: 5G enables the delivery of advanced driver-assistance features that reduce fatal accidents and traffic congestion. These features include real-time maps for navigation, speed warnings, road hazards, vulnerabilities, heads-up display systems, and sensor data sharing. These features will enable the vehicle to dynamically change its course on the road under certain scenarios and conditions. *Vehicle-to-Network* (V2N) communication is necessary for this use case. Information from the vehicle enables the remote application to perform short-range modelling and recognition of surrounding objects and vehicles plus mid- to long-range modelling of the surroundings using information on the latest digital maps, traffic signs, traffic-signal locations, road construction, and traffic congestion.
- *Autonomous Driving*: Fully autonomous driving involves the capability of a vehicle to sense its environment and navigate without human input under all scenarios and conditions. A 5G network with URLLC capability enables numerous necessary features, including the use of complex algorithms to distinguish between different cars on the road and identify appropriate navigation paths given obstacles and considering the rules of the road, and the exchange of information in real time between thousands of cars connected in the same area.
- *Tele-operated Driving*: This use case refers to the use of remote driver assistance in areas where automatic driving is not possible. This assistance can provide enhanced safety for disabled people, elderly populations, and drivers in complex traffic situations. Typical application scenarios include disaster areas and unexpected and difficult terrains for manual driving such as in mining and construction. Tele-operated driving requires the wireless network to support V2N communication of video, sound feed information, and diagnostics from the vehicle, along with environmental information, to the remote driver. The network must support transmitting control commands from the remote driver to the vehicle to maneuver the vehicle in real time.

## References and Further Reading

- [1] ITU-R, “IMT Vision—Framework and overall objectives of the future development of IMT for 2020 and beyond,” ITU-R Recommendation M.2083, September 2015.
- [2] ITU-R, “Minimum requirements related to technical performance for IMT-2020 radio interface(s),” ITU-R Report M.2410, November 2017.
- [3] ITU-T, “Requirements of the IMT-2020 network,” ITU-T Recommendation Y.3101, April 2018.
- [4] ITU-T, “High-level technical characteristics of network softwarization for IMT-2020,” ITU-T Recommendation Y.3150, January 2018.
- [5] ITU-T, “Framework of the IMT-2020 network,” ITU-T Recommendation Y.3102, May 2018.
- [6] 3GPP TS 22.261, “Technical Specification Group Services and System Aspects; Service requirements for the 5G system; Stage 1 (Release 18),” January 2021.
- [7] ITU-R, “Guidelines for evaluation of radio interface technologies for IMT-2020,” ITU-R Report M.2412, October 2017.
- [8] Mandar N. Kulkarni, Sarabjot Singh, and Jeffrey G. Andrew, “Coverage and Rate Trends in Dense Urban mmWave Cellular Networks,” 2014 *IEEE Global Communications Conference*, December 2014.
- [9] EU Project METIS, “Updated scenarios, requirements and KPIs for 5G mobile and wireless system with recommendations for future investigations,” Deliverable D1.5, April 2015.
- [10] Ericsson, “Cellular IoT Evolution for Industry Digitalization,” Ericsson White Paper, January 2019.
- [11] Ericsson, *Ericsson Mobility Report*, November 2020. (Issued annually).
- [12] Ericsson, “Cellular networks for Massive IoT,” Ericsson White Paper, January 2020.
- [13] ITU-T, “Overview of key performance indicators in smart sustainable cities,” ITU-T Recommendation Y.4900, June 2016.
- [14] ITU-T, “Supplement on use cases of smart cities and communities,” IT-T Series Y Supplement 56, December 2019.
- [15] 5G Americas, “New Services & Applications with 5G Ultra-Reliable Low-Latency Communications,” November 2018.
- [16] Next Generation Mobile Networks Alliance, “Verticals URLLC Use Cases and Requirements,” February 2020.
- [17] 5G Americas, “5G Services & Use Cases,” November 2017.



- [18] William Stallings, “Network Functions Virtualization,” *The Internet Protocol Journal*, Volume 24, No. 2, July 2021.
- [19] William Stallings, *5G Wireless: A Comprehensive Introduction*, ISBN-13: 9780136767299, Pearson, 2021.

WILLIAM STALLINGS is a consultant, lecturer, and author of over a dozen books on data communications and computer networking. He has a PhD in computer science from M.I.T. He has written numerous books on computer networking and computer architecture. His home in cyberspace is **WilliamStallings.com** and he can be reached at **ws@shore.net**

---

#### **Our Privacy Policy**

The *General Data Protection Regulation* (GDPR) is a regulation for data protection and privacy for all individual citizens of the *European Union* (EU) and the *European Economic Area* (EEA). Its implementation in May 2018 led many organizations worldwide to post or update privacy statements regarding how they handle information collected in the course of business. Such statements tend to be long and include carefully crafted legal language. We realize that we may need to provide similar language on our website and in the printed edition, but until such a statement has been developed here is an explanation of how we use any information you have supplied relating to your subscription:

- The mailing list for *The Internet Protocol Journal* (IPJ) is entirely “opt in.” We never have and never will use mailing lists from other organizations for any purpose.
- You may unsubscribe at any time using our online subscription system or by contacting us via e-mail. We will honor any request to remove your name and contact information from our database.
- We will use your contact information only to communicate with you about your subscription; for example, to inform you that a new issue is available, that your subscription needs to be renewed, or that your printed copy has been returned to us as undeliverable by the postal authorities.
- We will never use your contact information for any other purpose or provide the subscription list to any third party other than for the purpose of distributing IPJ by post or by electronic means.
- If you make a donation in support of the journal, your name will be listed on our website and in print unless you tell us otherwise.

# Low Earth Orbit Satellite Systems for Internet Access


by Dan York, Internet Society, and Geoff Huston, APNIC

Satellites have been providing Internet connectivity for a few decades now, but new technologies using satellites in *Low Earth Orbit* (LEO) have created a race to offer space-based Internet access that provides ubiquitous, high-speed, and low-latency connectivity. This article explores what is happening, what is new, and why our planet may soon be circled with 50,000 to 90,000 new satellites!

## Understanding Orbits

To understand the excitement over LEOs, you need to first understand the orbits of satellites (Figure 1).

Figure 1: Orbital Altitudes



## Geostationary Orbit

If you fire a projectile at a speed greater than the Earth’s “escape velocity,” (11.2 km/s) it will head away from the Earth. If you reduce the speed slightly, the projectile will be caught by the Earth’s gravity and try to fall back to Earth. If you also incline the aiming trajectory, then instead of falling back to Earth, it will settle into an orbit around the Earth. The orbital speed relative to the Earth is a function of the altitude of the object. At very high altitudes, such as the moon, the orbital period is slower than the Earth’s rotation, while at very low altitudes the orbital period is down to a small number of hours, implying that there is a mid-spot where the orbital period is the same as the Earth’s rotation. If you launch a satellite to an altitude of 35,768 km above the equator, orbiting in the same direction as the Earth on the equatorial plane, then from the Earth the spacecraft appears to sit in a stationary position when observed from the Earth’s surface; this orbit is a *Geostationary Orbit* (GEO)<sup>[1]</sup>. This geometry is what allows people to set up a satellite send/receive dish and point it at a specific location in the sky where the satellite is positioned—and never change the orientation of the dish.

GEO satellites are sufficiently distant that they can cover an entire hemisphere of the Earth's surface. However, they are normally equipped with a collection of transponders, most of which are focused on smaller areas, allowing the satellite to service multiple specific target regions at once, with greater total capacity as a result. A satellite operator can achieve global coverage with as few as three satellites. In addition to global service platforms, many nations have launched GEO satellites that are stationed over their country to provide communication services across their region.

These satellites are typically the size of a large bus and are expensive in terms of both construction costs and launch cost. Both the Moon and the Sun exert gravitational effects on the satellite, and, to a lesser extent, solar radiation pressure, all causing the satellite to drift away from its geosynchronous position. To counter this drift, the satellite is equipped with thrusters and some form of propellant. The total amount of onboard fuel defines an upper limit to the time that station position can be maintained, and these satellites typically have an operational lifespan of around 15–20 years.

In order to keep the level of radio interference between adjacent satellites to an acceptably low level, there are a limited number of geostationary orbit locations. Typically, geosynchronous satellite stations are separated by 2 degrees of angle as seen from the Earth, or 1,471 km apart in orbit. Disputes between nations over the deployment of satellites in this orbit are addressed through the coordination work hosted by the *International Telecommunications Union Radiocommunications Sector* (ITU-R).

A challenge with using GEO satellites for Internet access is that they are so far away from Earth. It takes a minimum of 238 ms for a signal to travel from the surface of the Earth to a satellite positioned 35,7683 km away and back again. The *Round Trip Time* (RTT) to propagate an outbound packet via a GEO satellite and receive a reply is a minimum of 476 ms. The distance to the satellite increases as you move away from the position directly underneath the satellite on the Earth's equator, and the propagation time for the round-trip approaches 560 ms as you approach the limit of clear signal access near the polar areas. When you add delays for signal encoding, switching, and other terrestrial elements, the delivered performance of a service based on geosynchronous satellites is a typical RTT of around 660 ms, or two-thirds of a second. For many applications that are tuned to operate efficiently on faster terrestrial paths, this extended delay often causes the application to be sluggish and unresponsive in terms of its performance.

It is also the case that at this altitude the Earth's magnetic field provides far less shielding from solar radiation via the *Van Allen Belt*, so the electronics for GEO satellites need to have appropriate shielding, and the onboard electronics must tolerate a certain amount of radiation exposure.

### Low Earth Orbit


An orbiting spacecraft needs to be positioned at least 160 km above the surface of the Earth, or it will encounter significant drag from the top of the Earth's atmosphere and its orbit will quickly decay, with an inevitable result. Above this altitude, it is viable to position orbiting spacecraft without needing to provide large quantities of continual propulsion (although some residual drag is experienced in orbital altitudes up to 500 km or so). For example, the International Space Station orbits at an altitude 400 km, with an orbital period of some 90 minutes. This region of space, where the orbits are higher than around 160 km and below 2,000 km, is termed the *Low Earth Orbit* (LEO) region. This is the region where we've positioned most of our satellites, as they are more accessible in terms of launch cost.

LEO satellites are close enough to the Earth's surface that signal propagation time to the satellite and back can be between 4 and 8 ms, which gives a range of RTT measurements for packet transmission via LEO services in the range 10–50 ms, a range comparable to that of terrestrial systems. With per-access service capacities of between 10 and 200 Mbps, LEO services can support most forms of modern real-time communication and online interaction<sup>[1]</sup>.

However, LEO-based satellite services require more complexity. At an altitude of 550 km, for example, a satellite will be visible from the Earth's surface in a circular area with a radius of some 900 km. Its orbital path is such that its velocity will be some 27,000 km/h, and each spacecraft will be visible from a fixed point on the Earth's surface for 5 minutes. In other words, to provide a continuous service over a fixed point, an evenly distributed collection of a minimum of 21 spacecraft would be needed in an orbital plane to ensure that as one satellite falls below the horizon, another is rising from the opposite horizon. Higher numbers of satellites in the orbital plane ensure a more reliable service and allow the Earth stations to avoid using spacecraft that are low in the horizon. To cover the entire Earth's surface, you need a minimum of 21 such orbital planes if you are using a 550-km altitude. The result is that, instead of just three satellites to provide a GEO service to anywhere on the planet, you would need hundreds or even thousands of satellites to provide the same comprehensive coverage. However, this scenario has some benefits in that the total capacity available from the system would also be many thousands of times greater in aggregate than the sparse GEO arrangement (refer to Figure 2).

LEO satellite systems may also be in multiple “shells” at different altitudes. For example, at the time of this article in July 2023, SpaceX's Starlink has launched 3,982 satellites into shells from 550 to 570 km, and another 750 satellites into shells from 525 to 535 km. Another provider, OneWeb, has launched 634 satellites into shells at 1,200 km in altitude<sup>[2]</sup>.

Figure 2: LEO Satellite Geometry




The closer a constellation is to the Earth's surface, the larger the number of satellites that are required for global coverage. OneWeb, orbiting at 1,200 km, needs only around 600 satellites, whereas SpaceX, orbiting around 550 km, needs around 3,000 satellites for continuous coverage. (They continue to launch more to provide additional capacity.) The counterpoint is that the closer satellites are to Earth, the lower the latency and the greater the potential access capacity.

LEO satellites are much smaller than GEO satellites and less expensive to manufacture and insert into orbit. For example, when SpaceX launches Starlink satellites into a LEO orbit, each Falcon 9 rocket usually carries 40 to 50 satellites. The development of reusable rockets has significantly affected the launch costs, bringing the launch cost on the SpaceX Falcon rockets to below USD 2,000 per kg. Other rocket systems have launch costs of between double to 20 times this cost.<sup>[3]</sup>

A challenge with LEO satellites is that because they are closer to the Earth, they are subject to periods of atmospheric drag when increased solar activity expands the extent of the upper atmosphere into the LEO orbital plane. Similar to GEO satellites, a LEO satellite must carry sufficient propellant to maintain its orbit and compensate for the gravitational influences of the Moon and the Sun that will otherwise pull it out of its orbital plane over time. Additionally, the operational lifespan of LEO satellites is typically around 5 years, so when it reaches its end of life it must have enough propellant in reserve to thrust the spacecraft down from its orbital plane ("de-orbit") and burn up in the upper atmosphere. When a LEO constellation is in active service, the system operator must employ a constant process of launching replacement satellites.

Tracking a LEO satellite using Earth equipment is challenging, because the tracking equipment must swing from horizon to horizon in approximately 5-minute intervals. Conventional parabolic dish antennae with mechanical pivots have been replaced by phased array antennae to allow the antenna unit to track satellites using electronic phase control. The ability to perform rapid phase shifts across the antenna array allows not only the tracking of these rapidly moving objects, but also rapid transitions from one satellite to the next to support connection handoffs between satellites (Figure 3).

Figure 3: Satellite Handoff



### Medium Earth Orbit

The region of space between 2,000 and 36,000 km above the Earth is referred to as *Medium Earth Orbit* (MEO). At the low end of this range, 2,000 km, the orbital period is around 2 hours, and the time to pass a signal up to the spacecraft and back is around 13 ms if you are located directly under the spacecraft. The higher the altitude the longer the orbital period, which, in turn reduces the number of spacecraft required to support a continuously available service. The transmission time to bounce a signal off the spacecraft at an altitude of 20,000 km is around 140 ms.

This zone encompasses the inner and outer Van Allen belts, which are belts of energetic charged particles that are trapped into an Earth orbit because of the Earth's magnetic field. The good news is that these belts protect the Earth's atmosphere from being blown away by solar radiation (as appears to have happened to Mars when its inner core solidified). The not-so-good news for satellites is that orbiting in this belt is like wandering through a firing range—there is always the possibility the sensitive electronics are damaged by a strike from one of these energetic particles. The outer belt is less dense, but the particles can have significantly higher energy levels. Beyond the Van Allen belts spacecraft encounter far higher levels of risk of damage from cosmic rays and solar radiation.

A region between the inner and outer Van Allen belts lies approximately between 12,000 and 24,000 km in altitude, which has a lower incidence of such energetic particles. The belts fluctuate in size and shape due to changes in the levels of solar radiation. The Earth itself acts as a shield, so that the belt is more compressed facing towards the Sun and extends further out on the “dark” side of the Earth.

The major satellites in this region are the satellite systems that support navigation, such as the *Global Positioning System* (GPS) using 31 spacecraft orbiting at some 20,200 km, *Galileo* with 24 active spacecraft at 23,222 km, *GLONASS* using 24 orbiting spacecraft at an altitude of 19,100 km, and *BeiDou* with 30 MEO satellites at 21,150 km.

For Internet access, the only major provider currently operating in MEO is the O3b network of around 20 satellites, operating at an altitude of approximately 8,000 km. O3b (which originally stood for “other 3 billion,” referring to the number of people still offline) began offering Internet connectivity in 2014 and was acquired by SES in 2016. SES is continuing to expand the service and is in the process of launching a new generation of 11 “O3b mPOWER” satellites into MEO, with promises to offer speeds up to multiple gigabits per second to its commercial customers.

Signal latency is higher than for LEOs, but significantly less than that of GEOs. Presumably feeling the competitive pressure from the LEO industry, SES has been working with many of the GEO providers to provide “multi-orbit” connectivity options that combine both MEO and GEO systems.

MEOs are a compromise in many ways. The Earth equipment still needs to perform tracking of the satellite, and that limits the power and sensitivity of the MEO antennae, yet the increased distance limits the performance and capacity of the system. The higher altitude provides greater coverage per satellite, allowing for broad coverage of the Earth’s surface with fewer spacecraft in the constellation, but the smaller number of satellites limits the overall capacity of the system.

If launch costs had remained high, then MEO systems made more sense in terms of minimizing the initial cost of the operation and maximizing the potential user base for the MEO satellite service, but the dramatic change in launch costs for LEO systems coupled with the use of phased array low-power steerable antennae has shifted the position quite dramatically in favor of LEO services. While GEO and MEO services tend to operate as wholesale services to a limited set of commercial customers using a conventional leased circuit service model, LEO systems have entered the consumer market, operating a direct access service as a retail service.

#### **User Equipment**

Consumers who want to connect to a LEO service for Internet access need to purchase what the satellite industry calls a *User Terminal*. This equipment includes the antenna and some access terminal, such as a small Ethernet switch or Wi-Fi access point. The phased-array antennas are compact, lightweight, and user-installable. Amazon has demonstrated some prototype user terminals that are small enough to fit in a backpack, and some mobile carriers have entered into agreements with Starlink to provide mobile handset access services directly from the handset to the satellite system (presumably with a significantly lower service capacity because of the limitations of the radio antenna on the handset).

There are more-complex user terminals that use parabolic dish antennae. These systems can achieve higher capacity and superior performance, but they require some form of mechanical steering to track the LEO satellite.



You can achieve the satellite-to-satellite handover function using dual antennae terminals, with one tracking the satellite for the active connection while the second pivots to focus on the next satellite in sequence.

#### **Earth Stations and Inter-Satellite Connections**

Satellite communications systems have conventionally operated as a “mirror in the sky.” The satellite receives a signal sent up from a user terminal and switches to a sending transponder that beams the signal back to an Earth station, which then passes the signal into the terrestrial network. With communications systems based on a GEO configuration it was possible to use just three Earth stations to service the entire system, given the hemispherical visibility of GEO satellites.

MEO and LEO satellites have more limited visibility, and in a “mirror” mode of operation, the density of Earth stations depends on the satellite altitude. For Starlink satellites at a 550-km altitude, Earth stations would need to be configured in a grid with around a 1,000-km spacing. This setup may be feasible in some locales, but if the intended service coverage includes more remote areas and coverage across oceans, then a different approach to Earth stations would be needed for LEO systems. An interesting technical development with LEO constellations to respond to this situation is the development of *Inter-Satellite Lasers* (ISLs) to send data between satellites within a constellation.


With the use of ISLs in the Starlink constellation, connections can now be passed from one satellite to another, and they do so in a relay form until they can drop down to an Earth station or user terminal. For example, SpaceX used this solution in early 2023 to connect remote Antarctic research stations, without having any ground stations on Antarctica<sup>[4]</sup>. Details of Starlink’s ISL system, including capacity and configuration, are not yet published, but there have been recent service announcements that show service availability in regions where no Earth stations exist.

A simple approach to ISLs would be to link adjacent satellites in a common orbital plane (Figure 4). A more flexible approach might be to link adjacent satellites across different orbital paths, but the relative closing speeds of satellites on different orbital planes may well be beyond the capabilities of steerable laser systems. For example, a spacecraft travelling north/south at 27,000 km/h attempting to track a satellite passing on an east-west plane at the same speed would have a closing speed of some 700 km/h and an angular velocity of one angle of degree every three seconds, which may present some challenges to an onboard laser steering system.

#### **Spectrum and Regulatory Approvals**

Satellite communication requires allocation of certain frequencies in the radio spectrum for communication from the consumer equipment (such as the antenna) to the satellites, and from the satellites to a ground station where connections are made to the Internet.

Figure 4: Inter-Satellite Lasers



Spectrum management could easily consume an entire article by itself, but at a high level, spectrum “allocations” are coordinated through the *International Telecommunication Union (ITU)*, specifically the *ITU Radiocommunication Sector (ITU-R)*. Most of the LEO systems use frequencies within the Ku (10.7–14.5 GHz) and Ka (17.3–30 GHz) bands, which are the bands that are intended for use by Broadband Satellite Services. There is also the more recently allocated Q/V band (37.5–51.4 GHz), which is available for use in this context and has been deployed already in some systems.

#### Initial Allocations

The world’s nations have agreed that above some altitude “outer space” began. The implication of this agreement is that the sovereign rights that apply to the defined surface parts of the Earth extend only up to the point of “outer space.” Oddly enough, the world’s nations did not agree as to where “outer space” begins, and some nations claim sovereignty up to an altitude of only 100 km, while others extend that further to 160 km or more. In any case, the result is that there is no national regime that must approve or otherwise say what a spacecraft may do in outer space in the form of “over flying” its territory. However, some conventions apply to assist various folk to coordinate their actions in space and assist in resolving any disputes that may arise. The use of GEO station slots by various nations is one such area where conventions apply, and the ITU-R assists in this coordination activity.

However, when the topic shifts to that of communications between Earth and satellites, there is a requirement to get various forms of national regulatory approval, based on the location of the Earth stations. The failure to gain such approvals for the Iridium service was the major cause of the early business failure of this venture in 2000. The approval is not quite as simple as just approval for the operation of Earth stations. When a company wants to launch a LEO satellite constellation, it conventionally obtains approval from its national regulator. For example, SpaceX and Amazon are both US companies, so they filed their requests with the US *Federal Communication Commission (FCC)*. The filings include the radio frequencies they want to use and the number, altitudes, and orbital planes at which their satellites will operate. These filings are also forwarded to the ITU.

The LEO and MEO space allocations are generally operated on a first-come, first-served basis, but to prevent people from “squatting” on spectrum and altitude allocations, the ITU requires LEO satellite constellations to have 10 percent of their constellation in orbit within the first two years after the start of deployment, 50 percent in five years, and 100 percent in seven years. This factor is part the reason why there is a great amount of heightened activity to launch LEO constellations. The various companies who have lodged applications need to meet these deployment milestones or they risk losing the exclusivity of their spectrum allocations.

For GEO satellites the ITU-R coordinates the spectrum allocations and the orbital slots. For LEO or MEO satellites, the ITU-R coordinates only the spectrum allocations, and does not coordinate the orbital planes and altitudes for these satellites. That aspect is handled entirely by national regulators.

When a company obtains the necessary spectrum and altitude allocations from its national regulator, it then can launch its satellite platforms into orbit. This launch activity requires completely separate approvals and involves processes whose descriptions fall outside the scope of this article.

#### **National Approvals**

As part of obtaining the initial spectrum allocations, a LEO system provider receives the approval to operate within its home nation. Then the provider has to go to the spectrum regulators in every country in which it wishes to operate to service and receive regulatory approval to use the spectrum in that country.

In some cases, the requested spectrum may be already in use, and the country faces the difficult decision of whether to re-configure its local use situation or be denied use of the LEO system. An example is Armenia, where a national regulator representative informed the audience of the *Armenian Internet Governance Forum 2022* that using Starlink was not possible anytime soon because the frequencies were in use by the Armenian military and government. Given that SpaceX will not change its frequencies, there probably will not be an option until the Armenian government changes its own systems to use different frequencies, which could take time and some amount of unplanned costs.

While it is technically possible for a LEO provider to offer service in a country for which it does not have permission to operate (SpaceX activated its service during the protests in Iran in late 2022<sup>[51]</sup>), it is not legally permitted to simply bring an antenna into a country and start using it with a LEO system. The LEO providers must obtain permission to operate the service in each country before they can make their service available to customers. Additionally, the LEO providers may also need to get approval to distribute the consumer equipment, and approval to interconnect with local terrestrial infrastructure.

### Standards

We have very little visibility into how the internal networks operate, but at the Internet Protocol and application layers, the LEO constellations so far seem to support all the conventional standards for Internet Protocol forwarding and Internet operations created by the *Internet Engineering Task Force* (IETF).

It seems at the moment that the major LEO operators (Starlink, OneWeb, Amazon Project Kuiper) are all pursuing their own proprietary systems for communication between their user terminals (antennas) and their satellites, and between their satellites and their ground stations. This process will require consumers to purchase completely separate user terminals in order to use each different system. Perhaps at some point this process will become standardized, but not in this initial period of deployment.

### Still Many Questions

That point is perhaps a critical one. Regardless of any marketing hype, the reality is that the LEO Internet access industry is very much still in its infancy. Only SpaceX's Starlink has global coverage. OneWeb has launched sufficient satellites to attain global coverage and is in the process of getting all its satellites in position. It hopes to offer global connectivity by the end of 2023, concentrating its service on the government and enterprise sector. Amazon's Project Kuiper has been manufacturing its satellites and equipment, but is still waiting for rocket availability to get its satellites into space. Many other companies are in various stages of getting their systems underway.

There are still many open questions, many of which the Internet Society explored in a recent document about LEO satellites<sup>[6]</sup>. What will the capacity of these LEO systems truly be? Will they be able to support all the many devices we want to connect to them? Will the systems be affordable by those who need the connectivity the most? Using space-based platforms to provide global coverage to the billions of unconnected people probably would require some significant changes in the service model, because the challenges, particularly in terms of affordability, are still significant.<sup>[8]</sup>

Will these constellations all be able to operate without interfering with each other? Will consumers tolerate the costs of proprietary equipment and the high cost of switching? What about the problems of "space debris" resulting from collisions or inactive satellites? Do we understand the potential environmental impact of having so many satellites burning up in our upper atmosphere when they reach their five-year end-of-life? Or the impact of all the regular rocket launches needed to resupply the constellations with new satellites? Questions abound, and many of them we may not be able to answer until we have the experience with getting more LEO constellations online.

### Looking Ahead

Without a doubt, the next few years will be extremely active:

- SpaceX plans to complete its “Gen1” constellation of 4,408 satellites. In addition, SpaceX has received US FCC approval for 7,500 satellites in its “Gen2” constellation, which the company hopes to grow to almost 30,000 satellites. The company has also announced numerous “direct-to-phone” services with a collection of mobile network operators.
- OneWeb aims to have its global connectivity service available by the end of 2023.
- Over the next two years, Amazon is seeking to launch its Project Kuiper, a direct competitor to Starlink in the consumer market, and have it operational in 2025.
- The Chinese government is seeking to launch its own LEO constellation, called “Guowang,” which will have almost 13,000 satellites.

Filings with the ITU show that there is a path where as many as 90,000 satellites could be launching into LEO over the next several years. Here are a few examples:

SpaceX Starlink Gen 1	4,408
SpaceX Starlink Gen 2	29,988
OneWeb, Phase 1	718
OneWeb, Phase 2	6,372
Amazon Project Kuiper	7,774
China Guowang	12,992
Astra	13,620
Boeing	5,842
Globalstar	3,080
Lynk	2,000
Telesat Lightspeed	1,969
Spin Launch	1,190
<b>TOTAL</b>	<b>89,953</b>

To put this information in perspective, LEO, MEO, and GEO orbits currently have only about 8,000 active satellites—and SpaceX operates over 4,500 of them!

Even more satellite systems are in planning stages; this article has covered only LEO satellites used for Internet access. In addition, LEO constellations are being launched for sensor networks (as in the Internet of Things), imaging/photography, and much more.

How many of these constellations will actually be launched into orbit and be used as a service is an open question. Navigating all the required regulatory approvals across all the various national regulatory regimes is very challenging. Unless you are SpaceX and own your own rockets, launching satellites into space is extremely difficult right now, as many current rocket programs are delayed or simply unavailable.

We are also seeing that the rise of the competition of space-based Internet systems is causing ground-based Internet service providers to accelerate their plans for deploying terrestrial networks. As good as LEO systems may be, fiber networks can still provide even higher speeds. There is still much in the way of potential to use fiber-based trunk networks with last-mile access provided by mobile radio technologies, and these networks are well-understood technologies with an already-dominant user base. If the terrestrial access service providers are successful in making even faster connectivity more widely available and affordable, then the competition for the user between space and terrestrial-based systems would increase in intensity, and we can anticipate that such competition will result in lower costs, wider coverage, and improved performance for consumers.

In any case, there are tremendous opportunities for LEO satellite systems to help us connect the unconnected, create more resilient networks, help coordinate disaster-relief efforts, and generally bring high-speed, low-latency Internet connectivity to everyone, everywhere. The next few years will show us whether we can make these goals a reality.

#### References and Further Reading

- [0] Note that in most satellite-related policies from the ITU and other regulators, geostationary orbit is abbreviated as “GSO,” and both LEO and MEO satellites are grouped as “Non-Geosynchronous Orbit” or “NGSO” satellites.
- [1] Josh Fomon, “New Speedtest Data Shows Starlink Users Love Their Provider,” *Ookla Insights Articles*, May 8, 2023.
- [2] Jonathan’s Space Pages:  
<https://planet4589.org/space/con/conlist.html>
- [3] *Our World in Data*, “Cost of space launches to low Earth orbit.”
- [4] Kevin Hurler, “Starlink Is Now Connecting Remote Antarctic Research Camps to the Internet,” *Gizmodo*, January 23, 2023.
- [5] Karl Vick, “Inside the Clandestine Efforts to Smuggle Starlink Internet Into Iran,” *TIME*, January 25, 2023.
- [6] “Perspectives on LEO Satellites: Using Low Earth Orbit Satellites for Internet Access,” Internet Society, 2022.
- [7] “sat-int” IETF mailing list for discussion of IETF technologies for satellite networking:  
<https://www.ietf.org/mailman/listinfo/Sat-int>

- [8] “Satellite broadband for the masses: Are we there yet?” Presentation by Mike Puchol to the APRICOT 2023/APNIC 55 Conference, March 2023.
- [9] Ulrich Speidel, “Getting hands-on experience with Starlink,” *APNIC Blog*, March 16, 2023.
- [10] Ulrich Speidel, “Everything you wanted to know about LEO satellites, Part 1: The basics,” *APNIC Blog*, May 20, 2021.
- [11] Ulrich Speidel, “Everything you wanted to know about LEO satellites, Part 2: Constellations, Gateways and Antennas,” *APNIC Blog*, May 27, 2021.
- [12] Ulrich Speidel, “Everything you wanted to know about LEO satellites, Part 3: Bandwidth, System Capacity and Inter-satellite Routing,” *APNIC Blog*, June 2, 2021.
- [13] Ulrich Speidel, “Everything you wanted to know about LEO satellites, Part 4: Why direct to site?,” *APNIC Blog*, June 11, 2021.
- [14] Ulrich Speidel, “Satellite still a necessity for many Pacific Islands,” *APNIC Blog*, September 18, 2018.

DAN YORK serves the *Internet Society* (ISOC) as the Director, Internet Technology, with his focus on helping explain the changes happening to the Internet and how we use it. He led ISOC’s 2022 *Low Earth Orbit* (LEO) satellite project and continues to monitor LEO systems as part of ISOC’s work to connect the unconnected. Dan is also active within the DNS and real-time communications areas of the *Internet Engineering Task Force* (IETF). Since the mid-1980s Dan has been working with online communication technologies and helping businesses and organizations understand how to use and participate in these new media. An author of multiple books on networking, security, IPv6, and Linux, he frequently presents at industry conferences and events. He has been blogging, writing online, and podcasting since the early 2000s. He lives in Vermont, USA. E-mail: [york@isoc.org](mailto:york@isoc.org)

GEOFF HUSTON AM, 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.net](mailto:gih@apnic.net)

---

### Check your Subscription Details!

Make sure that both your postal and e-mail addresses are up-to-date since these are the only methods by which we can contact you. If you see the words “Invalid E-mail” on your printed copy this means that we have been unable to contact you through the e-mail address on file. If this is the case, please contact us at [ipj@protocoljournal.org](mailto:ipj@protocoljournal.org) with your new information. The subscription portal is located here: <https://www.ipjsubscription.org/>



## Letter to the Editor

### On 25 years of The Internet Protocol Journal

Geoff Huston's magnum opus covering the last 25 years of the history of the Internet inspires me to reflect on those 25 years from several perspectives<sup>[1]</sup>. The year 1998 saw the launch of two extraordinary companies resting on and extending the Internet infrastructure: Google and Akamai. Each, in its own way, survived the "dot-bust" and went on to make significant contributions to the utility of the Internet and the World Wide Web. Many other companies have come, some have gone, and some have continued to enjoy robust growth. For many of us, 2023 feels like a new inflection point in communications. Low earth-orbiting satellites promise to provide access to the Internet from every square inch of the planet. Inter-continental fiber networks continue to expand in number and capacity, and there is energizing interest in developing mesh-like infrastructure to respond to cable cuts through optical re-routing in switching units on the ocean floor.

Ironically, 1998 was also the year in which the *Interplanetary Special Interest Group* of the Internet Society was formed, and it still continues as the Interplanetary chapter ([ipnsig.org](http://ipnsig.org)). A new suite of interplanetary "Bundle Protocols" has been standardized by the IETF and by the *Consultative Committee on Space Data Systems* (CCSDS), whose standards are also part of the ISO standards library. As the *Artemis*, *LunaNet* (NASA), and *MoonLight* (ESA) missions unfold, bringing us back to the Moon, the prospect of commercial operations looms large and immediate. The governance challenges of competitive and cooperative public/private engagements will be topics of urgent discussion well before this decade is out. There will be lessons from the multi-stakeholder policy-making practices derived from the building of the Internet and neo-institutional energy as the need for new governance mechanisms emerge.

Wireless technologies continue to evolve, and edge computing and more disciplined forms of Wi-Fi are emerging. Social media and recently emerging *Large Language Model* neural networks are confounding technologists and policy makers as they seek to make the Internet a safer place while preserving its historical openness to new ideas, new applications, and new ways to share and discover knowledge. These new neural transformer networks are living up to their name by transforming our awareness of the importance of reliable information sources in a sea of misinformation and disinformation. Preservation of human rights has become an ever more urgent priority in the face of scaled abuse of online resources. We can hope that the solutions to the problems of artificial intelligence will be found in the technology itself.

[1] Geoff Huston, "Twenty-Five Years Later," *The Internet Protocol Journal*, Volume 25, No. 1, June 2023.

And for all those years, *The Internet Protocol Journal* has consistently shed light on the conundrums that confound Internet engineers, scientists, operators, and policy makers. Hats off to its long-time editor, Ole Jacobsen, for persistent and quality reporting and sharing of timely and technically sound information.

—Vint Cerf, Woodhurst, May 2023

## Fragments

### Internet Society Launches NetLoss Calculator

The Internet Society recently launched the *NetLoss* calculator, a revolutionary tool that measures the economic impact of Internet shutdowns around the world. Hosted on the Internet Society's *Pulse* Platform that tracks and analyzes shutdowns, NetLoss uses a groundbreaking econometric framework to understand the impacts of shutdowns and provides an unprecedented level of rigor and precision in estimating their economic damage.

Internet shutdowns globally reached a record high in 2022, with governments around the world ordering Internet access and services to be restricted or blocked during civil unrest, school exams, and during elections, which resulted in major economic consequences.

According to the NetLoss calculator:

- The shutdown in Sudan in April 2023 is estimated to have cost the country more than \$3 million USD, as well as the loss of 560 jobs.
- The shutdown in Pakistan in May 2023 is estimated to have cost more than \$13 million USD, as well as increased unemployment.
- The shutdown in Guinea in June 2023 is estimated to have cost the country nearly half a million USD and job losses.

Governments often mistakenly believe that Internet shutdowns will quell unrest, stop the spread of misinformation, or reduce harm from cybersecurity threats. But shutdowns are extremely disruptive to economic activity: they halt e-commerce, generate losses in time-sensitive transactions, increase unemployment, interrupt business-customer communications, and create financial and reputational risks for companies. They also hurt a country's growth as research shows Internet adoption positively impacts *Gross Domestic Product* (GDP).

The Internet Society has long opposed the practice of Internet shutdowns, and urges all governments to refrain from implementing them due to the damage they inflict on a nation's economy, civil society, and Internet infrastructure. With NetLoss, organizations and advocates can demonstrate to governments and regulators how a shutdown will negatively impact their nation's economy.

In addition to the estimated cost of an Internet shutdown (that is, the loss in GDP), the Internet Society NetLoss calculator also estimates:

- The change in the unemployment rate due to a shutdown.
- The amount of *Foreign Direct Investment* (FDI) lost due to a shutdown.
- Risk of a shutdown: the probability that a country will experience a shutdown.

“The global rise in Internet shutdowns shows that governments continue to ignore the negative consequences of undermining the open, accessible, and secure nature of the global Internet. The calculator is a major step forward for the community of journalists, policymakers, technologists, and other stakeholders who are pushing back against the damaging practice of Internet shutdowns. Its groundbreaking and fully transparent methodology will help show governments around the world that shutting down the Internet is never a solution,” said Andrew Sullivan, President and CEO, of The Internet Society.

The calculator considers a wide range of economic impacts beyond traditional measures of economic output, such as GDP, to demonstrate the financial impact of an Internet shutdown. It also includes the change in the unemployment rate, the amount of FDI lost, and the risk of a future shutdowns.

In addition to its primary indicators, the NetLoss calculator’s methodology also takes into account other factors that can impact country-specific economic outcomes, including the age dependency ratio (percentage of working 18–65 years old to total population), the fraction of the population residing in urban areas, and the percentage of the labor force with basic education.

By using the following open data sets, the NetLoss calculator’s methodology is reproducible and transparent:

- *Shutdown Data*: Includes detailed event-level data on government-mandated shutdown events.
- *Protests and Civil Unrest*: Includes detailed event-level data on various events, their start and end dates, involved parties, and associated fatalities.
- *Elections*: The Constituency-Level Elections Archive maintained by Yale University provides elections data from 150 countries since 1960.
- *Socioeconomic Indicators*: The World Bank provides data on economic indicators including GDP per capita, employment, inflation, and foreign investment.

The framework used in the NetLoss calculator builds on the Internet Society’s longstanding research and advocacy on this issue via the Pulse Platform. Launched in December 2020, Internet Society Pulse consolidates trusted third-party Internet measurement data from various sources into a single platform to examine Internet trends and tell data-driven stories so that policymakers, researchers, journalists, network operators, civil society groups and others can better understand the health, availability, and evolution of the Internet. The source of data for NetLoss is the World Bank’s *World Development Indicators*, which typically corrects for minor statistical changes. Data used in the calculator is updated quarterly.

The NetLoss calculator can be found on the Pulse platform:  
**<https://pulse.internetsociety.org/netloss>**

### Reflections on Ten Years Past the Snowden Revelations

Authored by Stephen Farrell, Farzaneh Badii, Bruce Schneier, and Steven M. Bellovin, RFC 9446 contains the thoughts and recountings of events that transpired during and after the release of information about the United States *National Security Agency* (NSA) by Edward Snowden in 2013. There are four perspectives: that of someone who was involved with sifting through the information to responsibly inform the public, that of a security area director of the IETF, that of a human rights expert, and that of a computer science and affiliate law professor. The purpose of this memo is to provide some historical perspective, while at the same time offering a view as to what security and privacy challenges the technical community should consider. These essays do not represent a consensus view, but that of the individual authors. The RFC can be found here: <https://www.rfc-editor.org/info/rfc9446>

### APNIC Celebrates 30 Years

This year, the *Asia Pacific Network Information Centre* (APNIC) enters its third decade. Starting from a tiny office in Tokyo in 1992—with three people and a spreadsheet serving less than 600 delegated entities—APNIC has grown to a community of nearly 24,000 organizations across 56 economies. The APNIC of today serves 2.6 billion Internet users, more than half the global Internet. APNIC economies also comprise more than half the global IPv6 capability.

Despite many changes in technology and policy worldwide, APNIC has remained committed to: “A global, open, stable, and secure Internet.” The Asia Pacific is home to nine of the world’s 46 Least Developed Countries (as defined by the United Nations). Fifteen states and seven affiliated economies of APNIC Members are *Small Island Developing States* (SIDS), characterized by low population, distance and remoteness. However, the region also includes global economic superpowers and some of the most populated economies on Earth.

The combination of language, culture, distance, isolation and the different scale of the region’s communities magnifies the importance of consensus policy making on an equal basis. The APNIC community has developed policy that reflects and enables Internet growth across our region, and has ensured an Asia Pacific voice has been heard at a global level.

To mark this 30-year milestone, the *APNIC Blog* will run a series looking back at the past and into the future. The intention is to share stories, anecdotes, milestones and insights that capture some of the essence of the last 30 years. For more information visit:

<https://blog.apnic.net/2023/08/08/apnic-celebrates-30-years/>

## Thank You!

Publication of IPJ is made possible by 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. The following individuals have provided support to IPJ. You can join them by visiting <http://tinyurl.com/IPJ-donate>

Kjetil Aas	Ilija Bromberg	David Dillow	Serge Van Ginderachter	Richard Johnson
Fabrizio Accatino	Lukasz Bromirski	Richard Dodsworth	Greg Goddard	Jim Johnston
Michael Achola	Václav Brožík	Ernesto Doelling	Tiago Goncalves	Jonatan Jonasson
Martin Adkins	Christophe Brun	Michael Dolan	Ron Goodheart	Daniel Jones
Melchior Aelmans	Gareth Bryan	Eugene Doroniuk	Octavio Alfageme	Gary Jones
Christopher Affleck	Ron Buchalski	Michael Dragone	Gorostiaga	Jerry Jones
Scott Aitken	Paul Buchanan	Joshua Dreier	Barry Greene	Michael Jones
Jacobus Akkerhuis	Stefan Buckmann	Lutz Drink	Jeffrey Greene	Amar Joshi
Antonio Cuñat Alario	Caner Budakoglu	Aaron Dudek	Richard Gregor	Javier Juan
William Allaire	Darrell Budic	Dmitriy Dudko	Martijn Groenleer	David Jump
Nicola Altan	BugWorks	Andrew Dul	Geert Jan de Groot	Anders Marius Jørgensen
Shane Amante	Scott Burleigh	Joan Marc Riera	Ólafur Guðmundsson	Merike Kaeo
Marcelo do Amaral	Chad Burnham	Duocastella	Christopher Gumez	Andrew Kaiser
Matteo D'Ambrosio	Randy Bush	Pedro Duque	Gulf Coast Shots	Naoki Kambe
Selva Anandavel	Colin Butcher	Holger Durer	Sheryll de Guzman	Christos Karayiannis
Jens Andersson	Jon Harald Bøvre	Karlheinz Dölger	Rex Hale	Daniel Karrenberg
Danish Ansari	Olivier Cahagne	Mark Eanes	Jason Hall	David Kekar
Finn Arildsen	Antoine Camerlo	Andrew Edwards	James Hamilton	Stuart Kendrick
Tim Armstrong	Tracy Camp	Peter Robert Egli	Darow Han	Robert Kent
Richard Artes	Brian Candler	George Ehlers	Handy Networks LLC	Jithin Kesavan
Michael Aschwanden	Fabio Caneparo	Peter Eisses	Stephen Hanna	Jubal Kessler
David Atkins	Roberto Canonico	Torbjörn Eklöv	Martin Hannigan	Shan Ali Khan
Jac Backus	David Cardwell	Y Ertur	John Hardin	Nabeel Khatri
Jaime Badua	Richard Carrara	ERNW GmbH	David Harper	Dae Young Kim
Bent Bagger	John Cavanaugh	ESdatCo	Edward Hauser	William W. H. Kimandu
Eric Baker	Lj Cemerias	Steve Esquivel	David Hauweele	John King
Fred Baker	Dave Chapman	Jay Etchings	Marilyn Hay	Russell Kirk
Santosh Balagopalan	Stefanos Charchalakis	Mikhail Evstiounin	Headcrafts SRLS	Gary Klesk
William Baltas	Molly Cheam	Bill Fenner	Hidde van der Heide	Anthony Klopp
David Bandinelli	Greg Chisholm	Paul Ferguson	Johan Helsingius	Henry Kluge
A C Barber	David Chosrova	Ricardo Ferreira	Robert Hinden	Michael Kluk
Benjamin Barkin-Wilkins	Marcin Cieslak	Kent Fichtner	Damien Holloway	Andrew Koch
Feras Batainah	Lauris Cikovskis	Armin Fisslthaler	Alain Van Hoof	Ia Kochiashvili
Michael Bazarewsky	Brad Clark	Michael Fiumano	Edward Hotard	Carsten Koempfe
David Belson	Narelle Clark	The Flirble Organisation	Bill Huber	Richard Koene
Richard Bennett	Horst Clausen	Jean-Pierre Forcioli	Hagen Hultzs	Alexander Kogan
Matthew Best	James Cliver	Gary Ford	Kauto Huopio	Matthijs Koot
Hidde Beumer	Guido Coenders	Susan Forney	Asbjørn Højmark	Antonin Kral
Pier Paolo Biagi	Joseph Connolly	Christopher Forsyth	Kevin Iddles	Robert Krejčí
Arturo Bianchi	Steve Corbató	Andrew Fox	Mika Ilvesmaki	John Kristoff
John Bigrow	Brian Courtney	Craig Fox	Karsten Iwen	Terje Krogdahl
Orvar Ari Bjarnason	Beth and Steve Crocker	Fausto Franceschini	Joseph Jackson	Bobby Krupczak
Tyson Blanchard	Dave Crocker	Valerie Fronczak	David Jaffe	Murray Kucherawy
Axel Boeger	Kevin Croes	Tomislav Futivic	Ashford Jaggernaut	Warren Kumari
Keith Bogart	John Curran	Laurence Gagliani	Thomas Jalkanen	George Kuo
Mirko Bonadei	André Danthine	Edward Gallagher	Jozef Janitor	Dirk Kurfuerst
Roberto Bonalumi	Morgan Davis	Andrew Gallo	Martijn Jansen	Mathias Körber
Lolke Boonstra	Jeff Day	Chris Gamboni	John Jarvis	Darrell Lack
Julie Bottorff	Fernando Saldana Del	Xosé Bravo Garcia	Dennis Jennings	Andrew Lamb
Photography	Castillo	Oswaldo Gazzaniga	Edward Jennings	Richard Lamb
Gerry Boudreaux	Rodolfo Delgado-Bueno	Kevin Gee	Aart Jochem	Yan Landriault
Leen de Braal	Julien Dhallenne	Rodney Gehrke	Nils Johansson	Edwin Lang
Kevin Breit	Freek Dijkstra	Greg Giessow	Brian Johnson	Sig Lange
Thomas Bridge	Geert Van Dijk	John Gilbert	Curtis Johnson	Markus Langenmair



Fred Langham	David Millsom	Derrell Piper	Philip Schneck	Douglas Thompson
Tracy LaQuey Parker	Desiree Miloshevic	Rob Pirnie	James Schneider	Kerry Thompson
Alex Latzko	Joost van der Minnen	Jorge Ivan Pincay	Peter Schoo	Lorin J Thompson
Jose Antonio Lazaro	Thomas Mino	Ponce	Dan Schrenk	Fabrizio Tivano
Lazaro	Rob Minshall	Marc Vives Piza	Richard Schultz	Peter Tomsu Fine Art
Antonio Leding	Wijnand Modderman-	Victoria Poncini	Timothy Schwab	Photography
Rick van Leeuwen	Lenstra	Blahoslav Popela	Roger Schwartz	Joseph Toste
Simon Leinen	Mohammad Moghaddas	Andrew Potter	SeenThere	Rey Tucker
Robert Lewis	Charles Monson	Ian Potts	Scott Seifel	Sandro Tumini
Christian Liberale	Andrea Montefusco	Eduard Llull Pou	Paul Selkirk	Angelo Turetta
Martin Lillepuu	Fernando Montenegro	Tim Pozar	Andre Serralheiro	Michael Turzanski
Roger Lindholm	Roberto Montoya	David Raistrick	Yury Shefer	Phil Tweedie
Link Light Networks	Joel Moore	Priyan R Rajeevan	Yaron Sheffer	Steve Ulrich
Mike Lochocki	John More	Balaji Rajendran	Doron Shikmoni	Unitek Engineering AG
Chris and Janet Lonvick	Maurizio Moroni	Paul Rathbone	Tj Shumway	John Urbanek
Sergio Loreti	Brian Mort	William Rawlings	Jeffrey Sicuranza	Martin Urwaleck
Eric Louie	Soenke Mumm	Mujtiba Raza Rizvi	Thorsten Sideboard	Betsy Vanderpool
Adam Lovless	Tariq Mustafa	Bill Reid	Greipur Sigurdsson	Surendran Vangadasalam
Josh Lowe	Stuart Nadin	Petr Rejhon	Fillipe Cajaiba da Silva	Ramnath Vasudha
Guillermo a Loyola	Michel Nakhla	Robert Remenyi	Andrew Simmons	Randy Veasley
Hannes Lubich	Mazdak Rajabi Nasab	Rodrigo Ribeiro	Pradeep Singh	Philip Venables
Dan Lynch	Krishna Natarajan	Glenn Ricart	Henry Sinnreich	Buddy Venne
David MacDuffie	Naveen Nathan	Justin Richards	Geoff Sisson	Alejandro Vennera
Sanya Madan	Darryl Newman	Rafael Riera	John Sisson	Luca Ventura
Miroslav Madić	Thomas Nikolajsen	Mark Risinger	Helge Skrivervik	Scott Vermillion
Alexis Madriz	Paul Nikolich	Fernando Robayo	Terry Slattery	Tom Vest
Carl Malamud	Travis Northrup	Michael Roberts	Darren Sleeth	Peter Villemoes
Jonathan Maldonado	Marijana Novakovic	Gregory Robinson	Richard Smit	Vista Global Coaching
Michael Malik	David Oates	Ron Rockrohr	Bob Smith	& Consulting
Tarmo Mammers	Ovidiu Obersterescu	Carlos Rodrigues	Courtney Smith	Dario Vitali
Yogesh Mangar	Jim Oplotnik	Magnus Romedahl	Eric Smith	Rüdiger Volk
John Mann	Tim O'Brien	Lex Van Roon	Mark Smith	Jeffrey Wagner
Bill Manning	Mike O'Connor	Marshall Rose	Tim Sneddon	Don Wahl
Harold March	Mike O'Dell	Alessandra Rosi	Craig Snell	Michael L Wahrman
Vincent Marchand	John O'Neill	David Ross	Job Snijders	Lakhinder Walia
Normando Marcolongo	Carl Örne	William Ross	Ronald Solano	Laurence Walker
Gabriel Marroquin	Packet Consulting	Boudhayan	Asit Som	Randy Watts
David Martin	Limited	Roychowdhury	Ignacio Soto Campos	Andrew Webster
Jim Martin	Carlos Astor Araujo	Carlos Rubio	Evandro Sousa	Jd Wegner
Ruben Tripiana Martin	Palmeira	Rainer Rudigier	Peter Spekrijse	Tim Weil
Timothy Martin	Gordon Palmer	Timo Ruitter	Thayumanavan Sridhar	Westmoreland
Carles Mateu	Alexis Panagopoulos	RustedMusic	Paul Stancik	Engineering Inc.
Juan Jose Marin Martinez	Gaurav Panwar	Babak Saberi	Ralf Stempffer	Rick Wesson
Ioan Maxim	Chris Parker	George Sadowsky	Matthew Stenberg	Peter Whimp
David Mazel	Alex Parkinson	Scott Sandefur	Martin Štěpánek	Russ White
Miles McCredie	Craig Partridge	Sachin Sapkal	Adrian Stevens	Jurrien Wijlhuizen
Brian McCullough	Manuel Uruena Pascual	Arturas Satkovskis	Clinton Stevens	Derick Winkworth
Joe McEachern	Ricardo Patara	PS Saunders	John Streck	Pindar Wong
Alexander McKenzie	Dipesh Patel	Richard Savoy	Martin Streule	Makarand Yerawadekar
Jay McMaster	Dan Paynter	John Sayer	David Strom	Phillip Yaleloglou
Mark Mc Nicholas	Leif Eric Pedersen	Phil Scarr	Colin Strutt	Janko Zavernik
Olaf Mehlberg	Rui Sao Pedro	Gianpaolo Scassellati	Viktor Sudakov	Bernd Zeimetz
Carsten Melberg	Juan Pena	Elizabeth Scheid	Edward-W. Suor	Muhammad Ziad
Kevin Menezes	Chris Perkins	Jeroen Van	Vincent Surillo	Ziayuddin
Bart Jan Menkveld	Michael Petry	IngenSchenau	Terence Charles Sweetser	Tom Zingale
Sean Mentzer	Alexander Peuchert	Carsten Scherb	T2Group	Jose Zumalave
Eduard Metz	David Phelan	Ernest Schirmer	Roman Tarasov	Romeo Zwart
William Mills	Harald Pilz	Benson Schliesser	David Theese	廖明沂.

## 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](mailto:ole@protocoljournal.org) or [olejacobsen@me.com](mailto: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

<p><i>Supporters</i></p>  	<p><i>Diamond Sponsors</i></p> <p>Your logo here!</p>
<p><i>Ruby Sponsors</i></p> 	<p><i>Sapphire Sponsors</i></p> 

### *Emerald Sponsors*


















### *Corporate Subscriptions*









For more information about sponsorship, please contact [sponsor@protocoljournal.org](mailto:sponsor@protocoljournal.org)

---

The Internet Protocol Journal  
Link Fulfillment  
7650 Marathon Dr., Suite E  
Livermore, CA 94550

CHANGE SERVICE REQUESTED

---

## The Internet Protocol Journal

Ole J. Jacobsen, Editor and Publisher

### Editorial Advisory Board

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

**John Crain**, Senior Vice President and Chief Technology Officer  
Internet Corporation for Assigned Names and Numbers

**Dr. Steve Crocker**, CEO and Co-Founder  
Shinkuro, Inc.

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

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

**Dr. Cullen Jennings**, Cisco Fellow  
Cisco Systems, Inc.

**Olaf Kolkman**, Principal – Internet Technology, Policy, and Advocacy  
The Internet Society

**Dr. Jun Murai**, Founder, WIDE Project  
Distinguished Professor, Keio University  
Co-Director, Keio University Cyber Civilization Research Center, 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](mailto:ipj@protocoljournal.org)  
Web: [www.protocoljournal.org](http://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.*

