The Internet Protocol Journal

December 2023 Volume 26, Number 3

A Quarterly Technical Publication for Internet and Intranet Professionals

In This Issue

From the Editor 1

Introduction to 5G 2

Why ATM Failed..... 22

20 Years of Cellular and Wi-Fi Integration 30

RDRS 37

Fragments 39

Thank You! 40

Call for Papers 42

Supporters and Sponsors 43

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

ISSN 1944-1134

From The Editor

In our previous issue, we published Part One of "Introduction to 5G" by William Stallings. Part One introduced the standards and specifications that define 5G and described the usage scenarios that 5G supports. Part Two, included in this issue, provides an overview of the structure and function of 5G networks. A third article, on *Network Slicing*, which is closely related to 5G, will be published in a future edition of this journal.

This journal, as well as its predecessor *ConneXions—The Interoperability Report*, has covered numerous networking technologies over the last 35 years. Some of these technologies have become important building blocks for all networks, for example, *Ethernet*, which for more than 50 years has seen further improvements and standardization. We will publish an article on the history of Ethernet in a future issue. Other technologies have emerged, only to later fade into oblivion—an example being *Asynchronous Transfer Mode* (ATM). Our second article, by Craig Partridge, explores the reasons why ATM failed.

Modern smartphones and other mobile devices rely heavily on the use of numerous radio-based technologies such as *Near Field Communication* (NFC), Bluetooth, *Global Positioning System* (GPS), Wi-Fi, and cellular data. Our third article, by Mark Grayson, examines efforts to integrate cellular and Wi-Fi services into a single architecture.

The WHOIS protocol and its associated server were first introduced in 1982 in RFC 812. Described as "…a server ... that delivers the full name, U.S. mailing address, telephone number, and network mailbox for ARPANET users," the protocol specification was revised and finalized in RFC 3912 in 2004. WHOIS is an essential tool for anyone seeking information about a particular domain registration. Because of personal data protection laws, many ICANN-accredited registrars are now required to redact personal data from WHOIS lookups, yet certain parties may still have a legitimate need to access non-public information. ICANN has recently launched the *Registration Data Request Service* (RDRS) to address this need. Adiel Akplogan describes RDRS in our final article.

Publication of this journal is made possible by the generous support of our donors, supporters, and sponsors. In 2023 we were pleased to welcome *.au Domain Administration Limited* (auDA) and *Flexoptix* as our newest sponsors. If you would like to donate to or sponsor IPJ, please contact us at **ipj@protocoljournal.org**

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

Introduction to 5G

Part Two: Core Network, Radio Access Network, and Air Interface

by William Stallings

Part One of this 5G introduction[0] addressed requirements, standards, and applications for 5G. This part provides an overview of the structure and function of 5G networks.

Core Network

Dozens of 3GPP specifications are related to the 5G core network that together describe a system of significant complexity. A key document in this collection is *3GPP Technical Specification TS 23.501*[1]. This document, 450 pages long in its current release, provides a detailed technical overview of the core network architecture, procedures, services, and interfaces, and is the primary basis for the presentation in this article.

The core network architecture can be viewed as a set of interconnected *Network Functions* (NFs). An NF is a processing function in a network, which has defined functional behavior and interfaces. An NF can be implemented as a network element on dedicated hardware, a software instance running on dedicated hardware, or a virtualized function instantiated on an appropriate platform.

Figure 1: 5G Core Network Functional Architecture

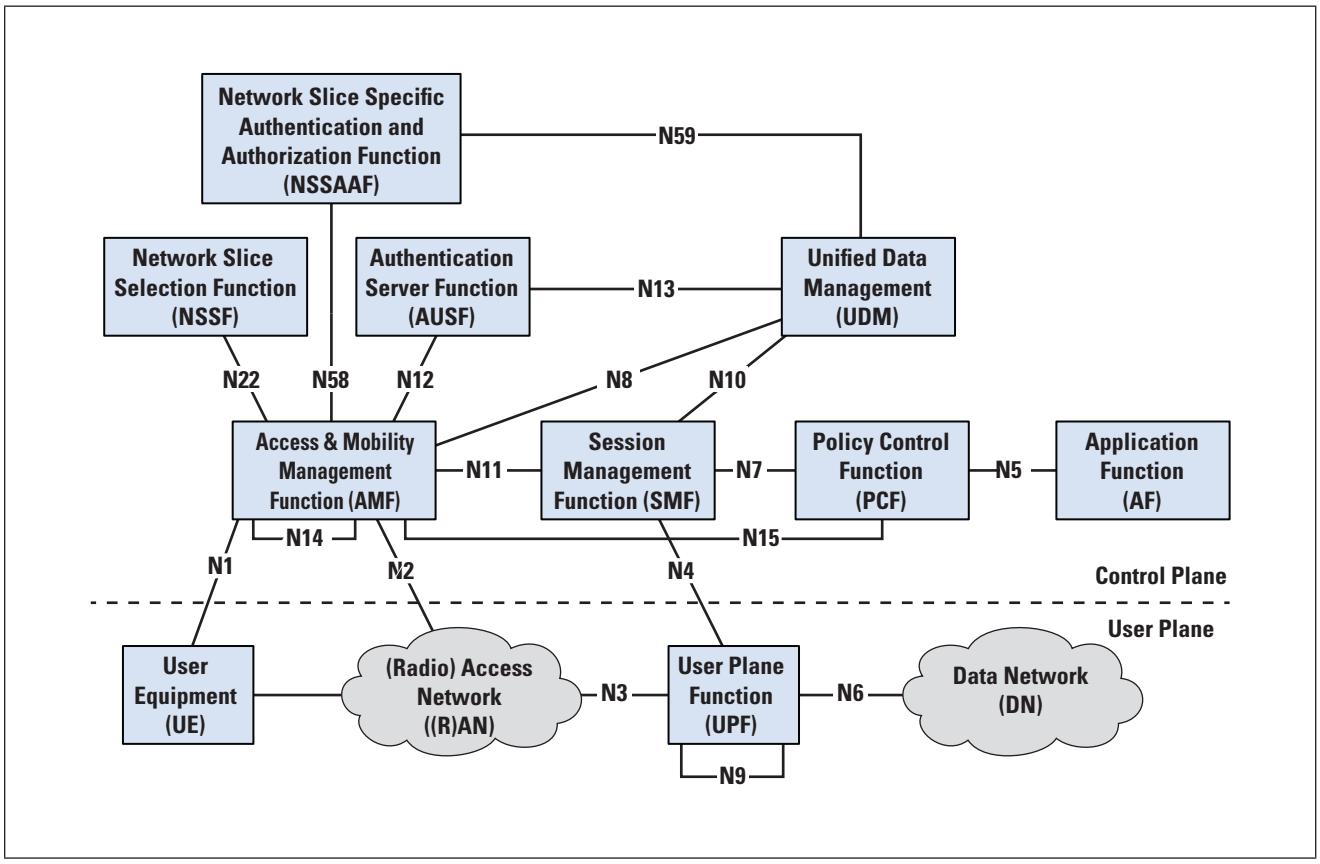

TS 23.501 contains numerous architecture diagrams from several different points of view and at varying levels of detail. Figure 1 depicts the basic 5G architecture using a reference-point representation, showing how the NFs interact with each other.

The figure includes the following NFs and other modules:

- *Authentication Server Function* (AUSF): Performs authentication between *User Equipment* (UE) and the network
- *Access and Mobility Management Function* (AMF): Receives all connection- and session-related information from the UE (N1/N2) but is responsible only for handling connection, registration, reachability, and mobility management tasks. All messages related to session management are forwarded to the *Session Management Function* (SMF).
- *Network Exposure Function* (NEF): Provides an interface for outside applications to communicate with the 5G network to obtain network-related information about the capabilities of the network.
- *Network Repository Function* (NRF): Allows NFs to register their functionality and to discover the services offered by other NFs present in the network.
- *Network Slice Selection Function* (NSSF): Selects the set of network slice instances to accommodate the service request from a UE. When a UE requests registration with the network, AMF 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.
- *Network Slice-Specific Authentication and Authorization* (NSSAAF): Performs authentication and authorization specific to a slice.
- *Policy Control Function* (PCF): Provides functionalities for the control and management of policy rules including rules for *Quality of Service* (QoS) enforcement, charging, and traffic routing. PCF enables end-to-end QoS enforcement with QoS parameters (for example, maximum bit rate, guaranteed bit rate, and priority level) at the appropriate granularity (for example, per UE, per flow, and per PDU session.
- *Session Management Function* (SMF): Responsible for *Protocol Data Unit* (PDU) session establishment, modification, and release between a UE and a data network. A PDU session, or simply session, is an association between the UE and a data network that provides a PDU connectivity service. A PDU connectivity service is a service that provides for the exchange of PDUs between a UE and a data network.
- *Unified Data Management* (UDM): Responsible for access authorization and subscription management. UDM works with AMF and AUSF as follows: The AMF provides UE authentication, authorization, and mobility management services. The AUSF stores data for authentication of UEs, and the UDM stores UE subscription data.

- *User Plane Function* (UPF): Handles the user plane path of PDU sessions. UPF functions include packet routing and forwarding, QoS handling, traffic usage reporting, and policy rule enforcement.
- *Application Function* (AF): Provides session-related information to PCF so that SMF can ultimately use this information for session management. 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.
- *User Equipment* (UE): Gives users access to network services. An example is a mobile phone. For the purpose of 3GPP specifications, the interface between the UE and the network is the radio interface.
- *(Radio) Access Network* [(R)AN]: A network that provides access to a 5G core network. It includes the 5G RAN and other wireless and wired access networks.
- *Data Network* (DN): A network to which UE is logically connected by a session. It may be the Internet, a corporate intranet, or an internal services function within the mobile network operator's core (including content-distribution networks).
- *Service Communication Proxy* (SCP): NFs and NF services can communicate directly or indirectly via the SCP. The SCP enables multiple NFs to communicate with each other and with user plane entities in a highly distributed multi-access edge compute cloud environment. These services provide routing control, resiliency, and observability to the core network.

In Figure 1, two reference points loop back to the same function: N9 and N14. The N9 reference point is an interface between two distinct UPFs used for forwarding packets. The N14 reference point is between two AMFs, one acting as a source AMF for a data transfer and the other acting as a destination AMF.

An example of the interaction of the various NFs is the session-establishment procedure, which is defined in TS 23.502[2]. Figure 2 provides a much-simplified view of the interaction between the various network components during session establishment. Session establishment begins with a request from the UE over the RAN, which is directed to the AMF. An SMF is selected to manage the PDU session. SMF uses UDM in the process of creating a session and performing authentication and authorization. SMF selects a PCF for the session and a UPF to handle data plane PDU forwarding in both directions. SMF establishes a session with the DN. After a few more exchanges, the UE is able to communicate over a session with the DN.

SDN and NFV

Two essential enablers of 5G services provided by core networks are *Software-Defined Networking* (SDN) and *Network Functions Virtualization* (NFV).[15]

ITU-T Y.3300[3] defines SDN as a set of techniques that enables users to directly program, orchestrate, control, and manage network resources, thereby facilitating the design, delivery, and operation of network services in a dynamic and scalable manner.

Figure 2: UE-Requested PDU Session Establishment

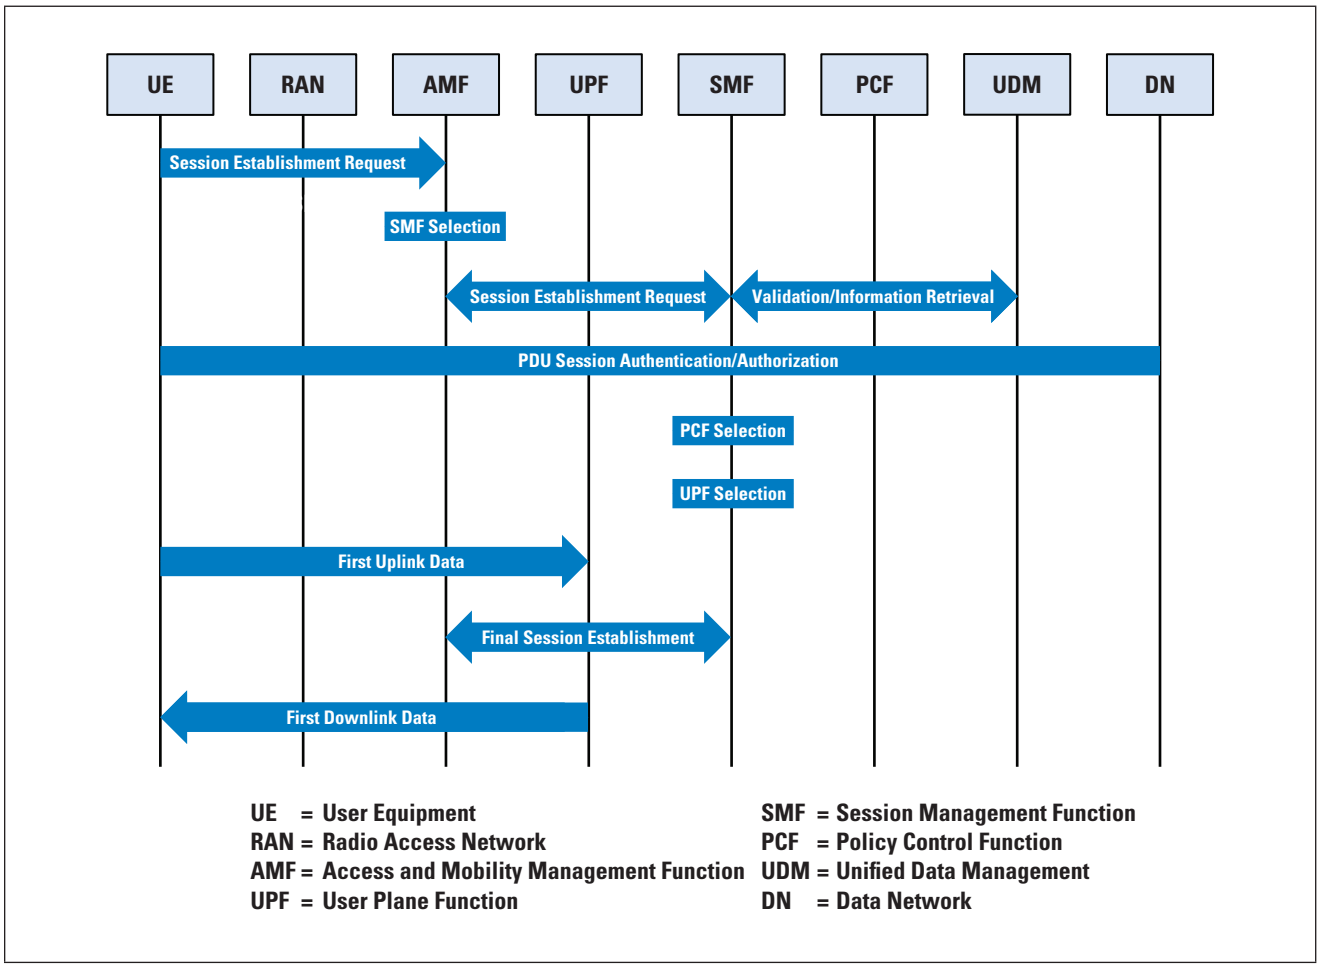

The two elements involved in forwarding packets through routers are a *control function*, which decides the route the traffic takes and the relative priority of traffic, and a *data function*, which forwards data based on control-function policy. Prior to SDN, these functions were performed in an integrated fashion at each network device (router, bridge, packet switch, etc.). Control in such a traditional network is exercised with a routing and control network protocol that is implemented in each network node. This approach is relatively inflexible and requires all of the network nodes to implement the same protocols. With SDN, a central controller performs all complex functionality, including routing, naming, policy declaration, and security checks. This central controller constitutes the SDN control plane, and consists of one or more SDN controllers. The SDN controller defines the data flows that occur in the SDN data plane. Each flow through the network is configured by the controller, which verifies that the communication is permissible by the network policy.

If the controller allows a flow requested by an end system, it computes a route for the flow to take, and adds an entry for that flow in each of the switches along the path. With all complex functions subsumed by the controller, switches simply manage flow tables whose entries can be populated only by the controller. The switches constitute the data plane. Communication between the controller and the switches uses a standardized protocol.

Figure 3 illustrates the SDN architecture. The data plane consists of physical switches and virtual switches, both of which are responsible for forwarding packets. The internal implementation of buffers, priority parameters, and other data structures related to forwarding can be vendor-dependent. However, each switch must implement a model, or an abstraction, of packet forwarding that is uniform and open to the SDN controllers. This model is defined in terms of an open *Application Programming Interface* (API) between the control plane and the data plane (that is, the southbound API).

Figure 3: SDN Architecture

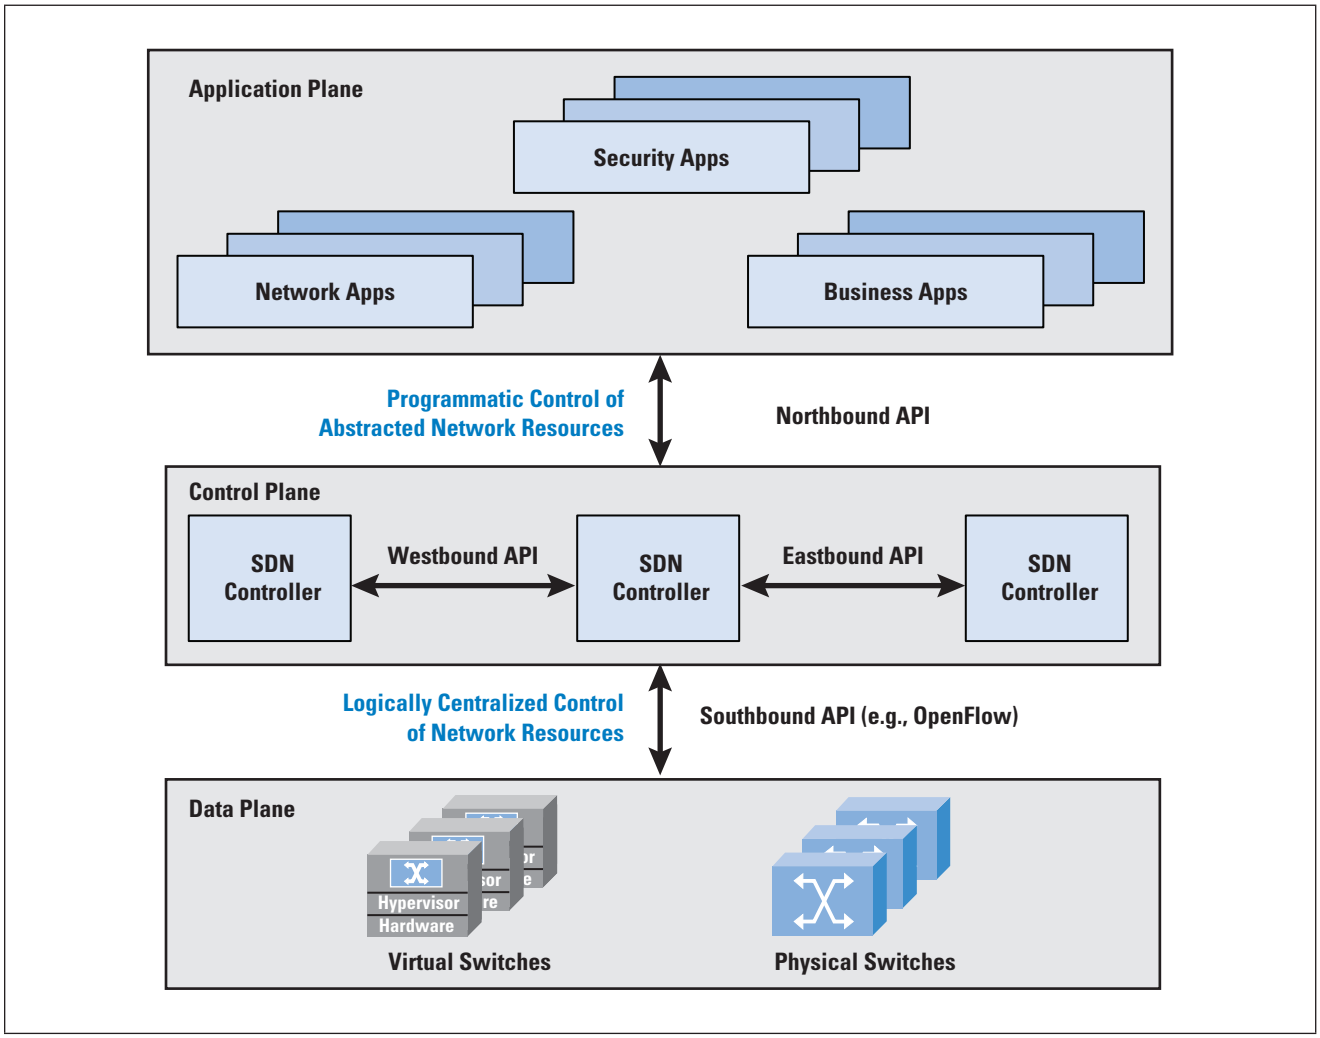

Similarly, SDN controllers can be implemented directly on a server or on a virtual server. An API is used to control the switches in the data plane. In addition, controllers use information about capacity and demand obtained from the networking equipment through which the traffic flows. SDN controllers also expose northbound APIs, meaning that developers and network managers can deploy a wide range of off-the-shelf and custom-built network applications, many of which were not feasible prior to the advent of SDN.

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.

Figure 4 shows a high-level view of the NFV framework, which supports the implementation of network functions as software-only NVFs.

Figure 4: High-Level NFV Framework

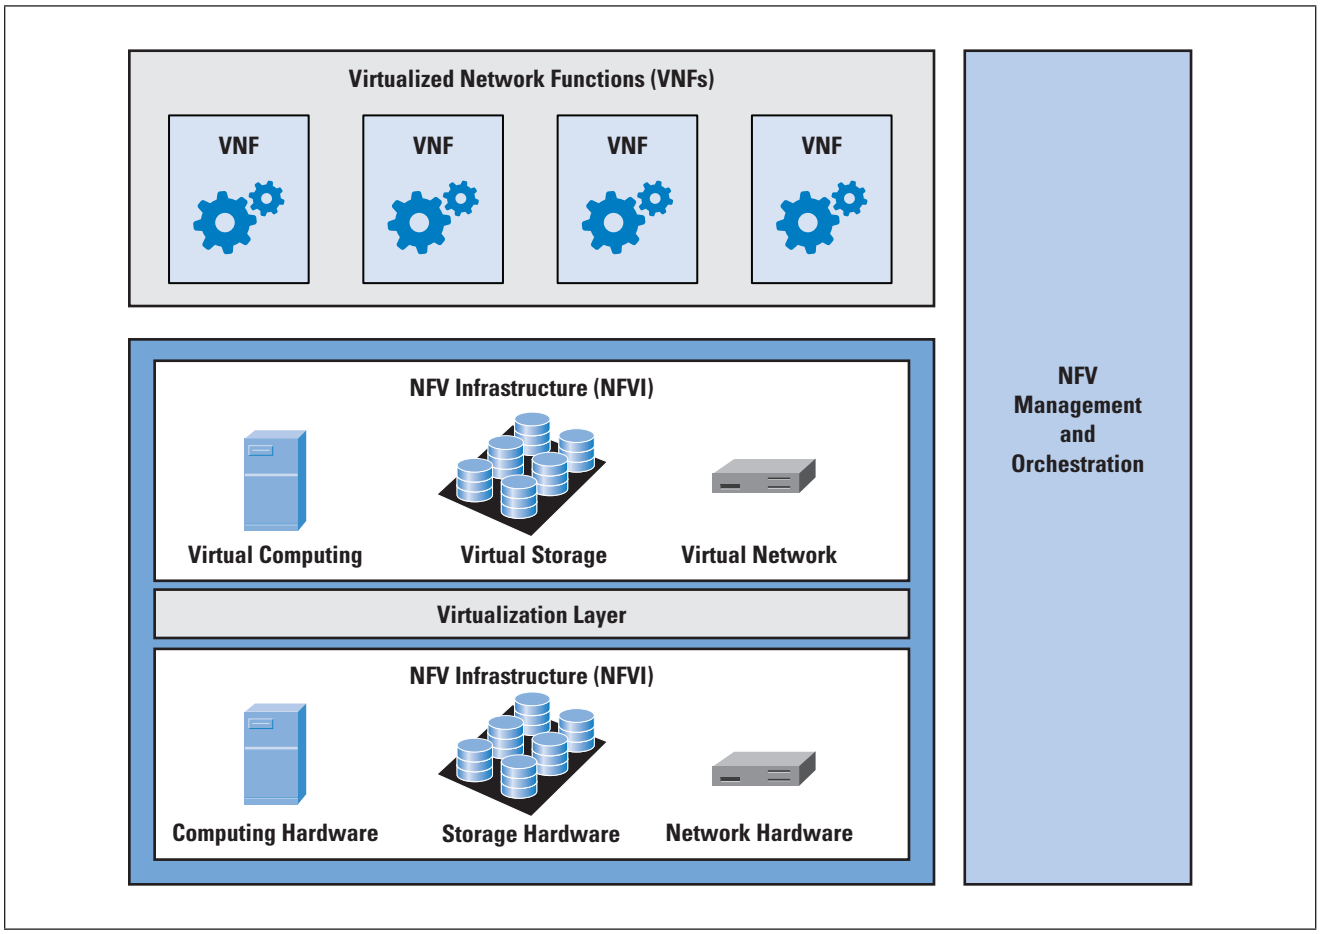

The NFV framework consists of three domains of operation:

- *Virtualized Network Functions* (VNFs): These functions consist of a collection of VNFs, implemented in software, run over the *NFV Infrastructure* (NFVI)
- *NFV Infrastructure* (NFVI): The NFVI performs a virtualization function on the three main categories of devices in the network service environment: computer devices, storage devices, and network devices.
- *NFV Management and Orchestration:* This domain encompasses the orchestration and lifecycle management of physical and/or software resources that support the infrastructure virtualization and the lifecycle management of VNFs. NFV management and orchestration focuses on all virtualization-specific management tasks necessary in the NFV framework.

NFV and SDN are independent but complementary schemes. SDN decouples the data and control planes of network traffic control, making the control and routing of network traffic more flexible and efficient. NFV decouples network functions from specific hardware platforms via virtualization to make the provision of these functions more efficient and flexible. Virtualization can be applied to the data plane functions of the routers and other network functions, including SDN controller functions. Thus, either can be used alone, but the two can be combined to reap greater benefits.

The European Telecommunications Standards Institute (ETSI) document GS NFV-EVE 005[4] examines the manner in which SDN can be incorporated in the NFVI to provide connectivity services.

The framework incorporates two controllers: one logically placed at the tenant level and another at the NFVI level. Each controller centralizes the control plane functionalities and provides an abstract view of all the connectivity-related components it manages. The controllers are as follows:

- *Infrastructure SDN Controller* (IC): This controller enables communication among VNFs and among their components, including the cases when those VNFs are instantiated in separated *Points of Presence* (PoPs), reachable through a WAN connection. Managed by the *Virtualized Infrastructure Manager* (VIM), this controller may change infrastructure behavior on demand according to VIM specifications adapted from tenant requests.
- *Tenant SDN Controller* (TC): Instantiated in the tenant domain as one of the VNFs or as part of the *Network Management System* (NMS), this second controller dynamically manages the pertinent VNFs used to realize the tenant's network service(s). These VNFs are the underlying forwarding plane resources of the TC. The operation and management tasks that the TC carries out are triggered by the applications running on top of it (for example, the *Operations Support System* [OSS]).

Figure 5: Integrating SDN Controllers into the Reference NFV Architectural Framework

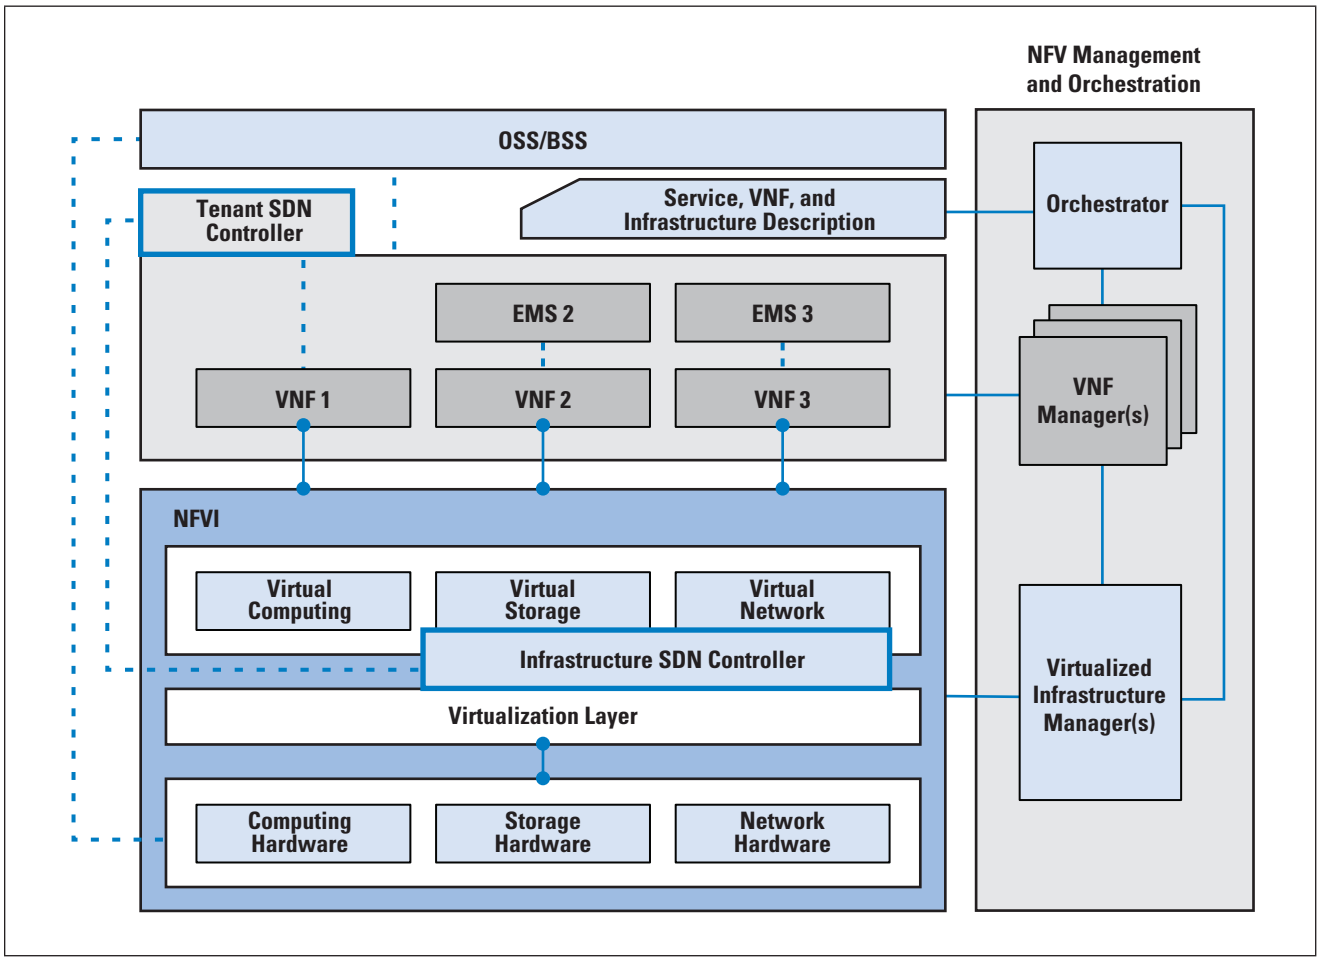

Network Slicing

One of the most important features of 5G is *Network Slicing*. Indeed, network slicing is essential to the exploitation of the capabilities defined for 5G. Network slicing enables a 5G network operator to provide customized networks by creating multiple virtual and end-to-end networks, referred to as network slices. Each network slice can be defined according to different requirements on functionality, QoS, and specific users.

"Network Slicing for 5G: Challenges and Opportunities"[5] lists the following advantages of slicebased networking compared with traditional networking:

- Network slicing can provide logical networks with better performance than one-size-fits-all networks.
- A network slice can scale up or down as service requirements and the number of users change.
- Network slices can isolate the network resources of one service from the others; the configurations among various slices don't affect each other. Therefore, the reliability and security of each slice can be enhanced.
- A network slice is customized according to QoS requirements, which can optimize the allocation and use of physical network resources.

Network slicing is made possible by NFV and SDN. NFV implements the *Network Functions* (NFs) in a network slice, enabling the isolation of each network slice from all other network slices. Isolation can be achieved by one or more of the following: (1) using a different physical resource; (2) separation by virtualization, which may allow sharing of physical resources; or (3) through sharing a resource with the guidance of a respective policy that defines the access rights for each tenant. Isolation assures QoS and security requirements for that slice, independent of other slices operating on the network from the same or different users. After a network slice is defined, SDN operates to monitor and enforce QoS requirements by controlling the behavior of the QoS flow for each slice.

Figure 6, based on concepts in a *Next Generation Mobile Networks* report[6], illustrates network slicing concepts.

The figure shows a simple core network configuration composed of three types of devices:

1. *Cloud Nodes:* These nodes provide cloud services, software, and storage resources. There are likely to be one or more central cloud nodes that provide traditional cloud computing service. In addition, cloud-edge nodes provide low-latency and higher-security access to client devices at the edge of the network. All of these nodes include virtualization system software to support virtual machines and containers. NFV enables effective deployment of cloud resources to the appropriate edge node for a given application and given fixed or mobile user. The combination of SDN and NFV enables the movement of edge resources and services to dynamically accommodate mobile users.
2. *Networking Nodes*: These nodes are IP routers and other types of switches for implementing a physical path through the network for a 5G connection. SDN provides for flexible and dynamic creation and management of these paths.
3. *Access Nodes:* These nodes provide an interface to radio access networks, which in turn provide access to mobile UE. SDN creates paths that use an access node for one or both ends of a connection involving a wireless device.

The remainder of Figure 6 illustrates three use cases. The blacked-out core network resources represent resources not used to create the network slice. Cloud nodes that are part of the slice may include the following:

- Control plane functions associated with one or more user plane functions (for example, a reusable or common framework of control)
- Service- or service category-specific control plane and user plane function pairs (for example, a user-specific multimedia application session)

Figure 6: 5G Network Slices Implemented on the Same Infrastructure

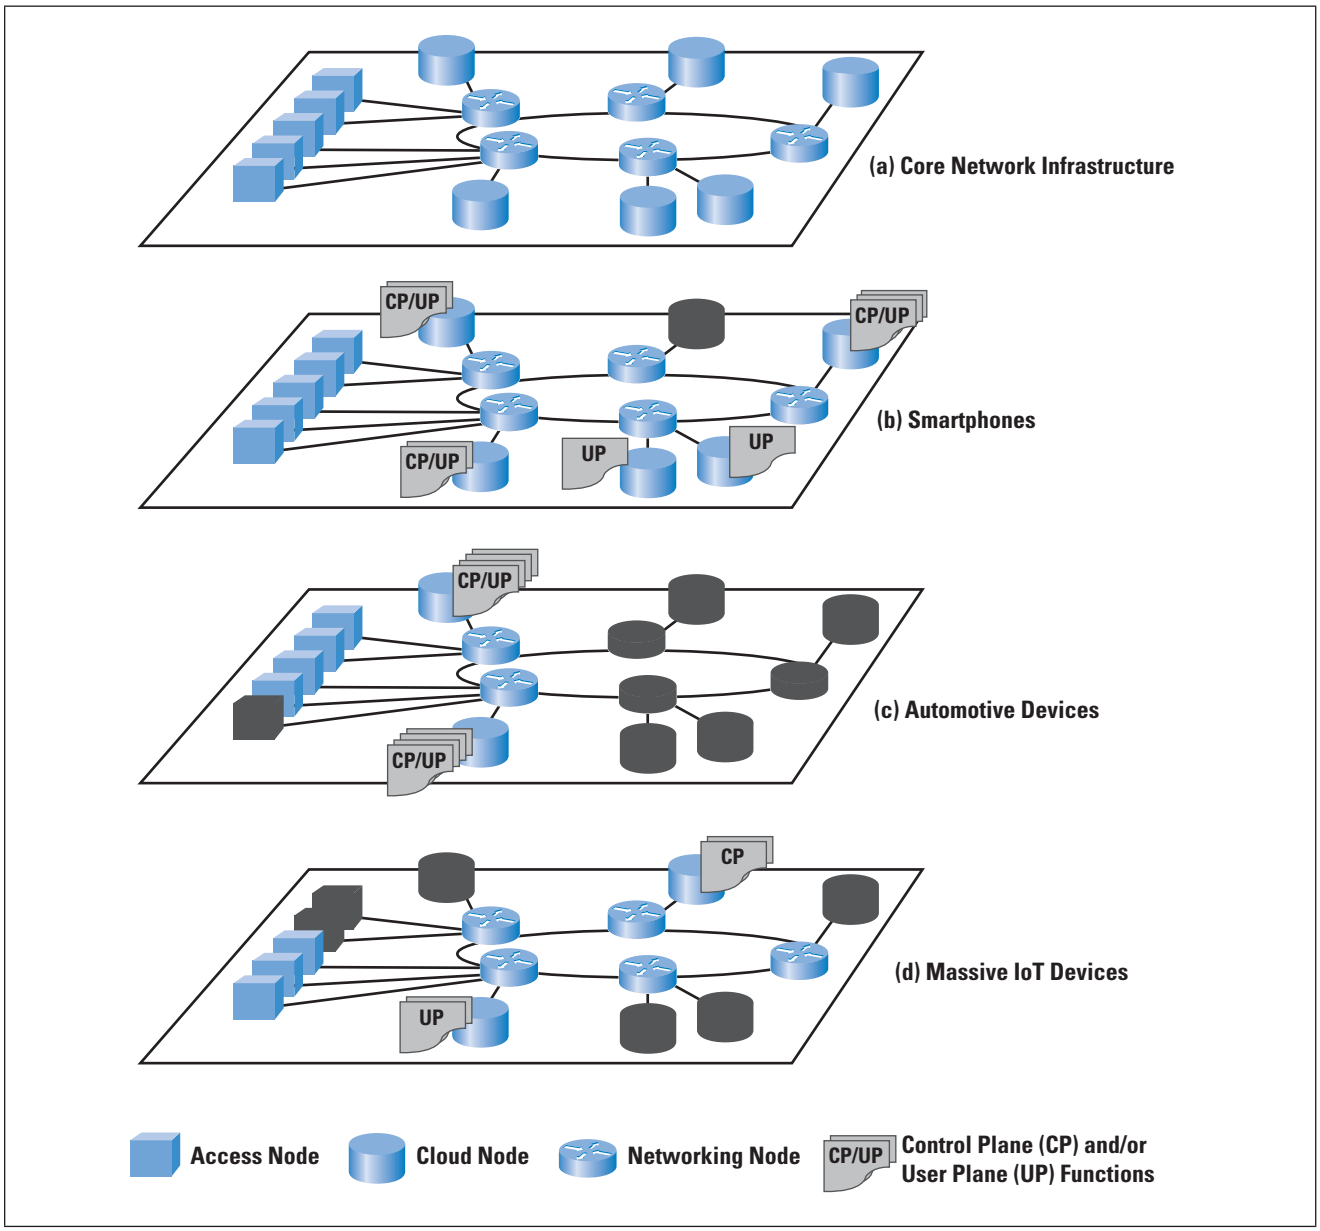

The first network slice depicted in Figure 6 is for a typical smartphone use case. Such a slice might have fully fledged functions distributed across the network. The second network slice in the figure indicates the type of support that may be allocated for automobiles in motion. This use case emphasizes the need for security, reliability, and low latency. A configuration to achieve these needs would limit core network resources to nearby cloud edge nodes, plus the recruitment of sufficient access nodes to support the use case. The final use case illustrated in Figure 6 is for a massive *Internet of Things* (IoT) deployment, such as a huge number of sensors. The slice can contain just some specific *Control Plane* (CP) and *User Plane* (UP) functions with, for example, no mobility functions. The CP and UP functions might include filtering and preliminary data analysis at the edge and big data types of analysis at a more central node. This slice would need to engage only access nodes nearest to the IoT device deployment.

Radio Access Network

As for the 5G core network, there are dozens of 3GPP specifications related to the 5G RAN. A key document in this collection is *3GPP Technical Specification TS 38.300*[7]. This document provides a detailed technical overview of the RAN architecture, protocols, functionality, and interfaces, and is the primary basis for the presentation in this article.

The overall RAN architecture, in terms of the deployment of base-station RAN nodes, is dictated by the need to coexist with 4G UE and 4G core networks for an extended period. In 2019, 4G became the dominant mobile technology across the world, with more than 4 billion connections, accounting for 52% of total connections (excluding licensed cellular IoT). 4G connections will continue to grow for the next few years, peaking at just under 60% of global connections by 2023[8]. It is clear that 4G UE will form a major portion of the cellular demand for quite a few years to come.

The most important requirement for 5G carriers is to provide full support for both 4G and 5G UE. Figure 7, from TS 38.300, is a simplified view of the overall RAN architecture and its interface to the 5G core network for providing that support. The figure depicts two types of base stations. The gNB node provides *5G New Radio* (NR) user plane and control plane protocol terminations toward the UE and connects via the NG interface to the *5GC* (5G core). The ng-eNB node provides 4G, referred to as *Evolved Universal Terrestrial Radio Access* (E-UTRA) user plane and control plane protocol terminations toward the UE and connects via the NG interface to the 5GC.

Figure 7: Overall Radio Access Network Architecture

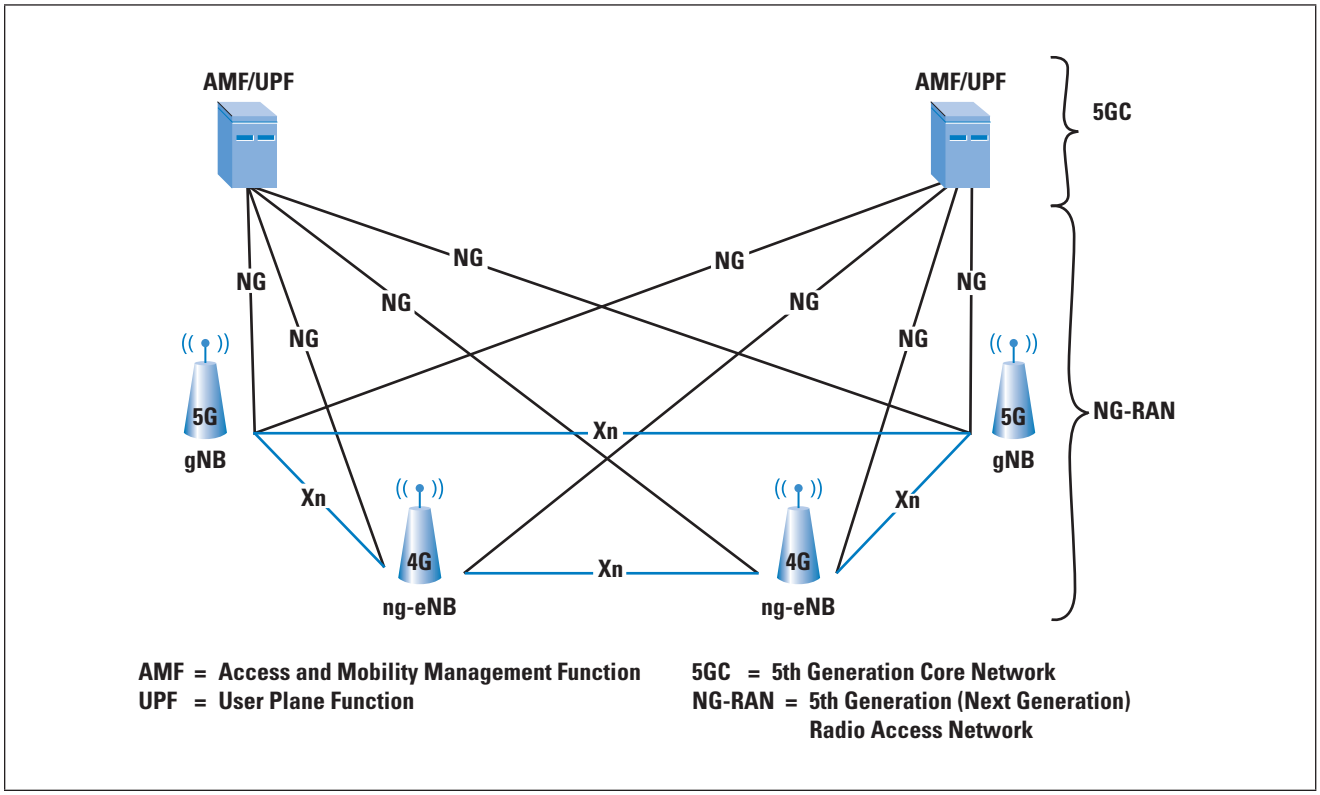

The base stations interconnect with each other by means of the Xn interface. Base stations connect to the core network through the NG interfaces, more specifically to the AMF (access and mobility management function) by means of the NG-C interface and to the UPF by means of the NG-U interface.

Figure 8, from TS 38.401[9], provides a different perspective on key 5G RAN interfaces. A gNB may be a single integrated system, referred to as a *monolithic* or *non-split* node. Or a gNB may be organized as a split node, consisting of a *gNB-Central Unit* (gNB-CU) and one or more *gNB-Distributed Units* (gNB-DUs). The CU processes non-real-time protocols and services, and the DU processes physical layer protocol and real-time services. One gNB-DU supports one or multiple cells. One cell is supported by only one gNB-DU. A gNB-CU and the gNB-DU units are connected via the F1 logical interface. One gNBDU is connected to only one gNB-CU. For resiliency, a gNB-DU may also be able to connect to another gNB-CU (if the primary gNB-CU fails) through appropriate implementation. NG, Xn, and F1 are logical interfaces.

Figure 8: RAN Interfaces

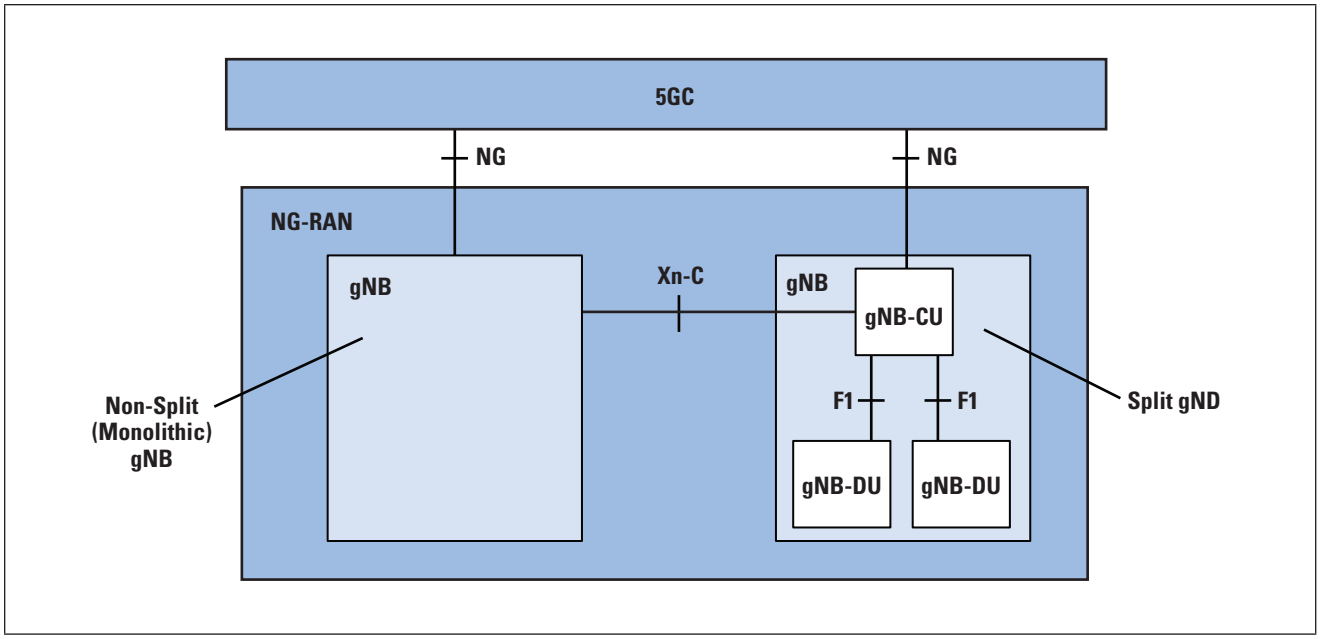

Air Interface

As with other aspects of 5G, dozens of 3GPP specifications are related to the 5G RAN. However, the definitive document is *ITU-R Recommendation M.2150*[10], issued in February 2021. The current version of the Recommendation adopts three radio interface technologies: 3GPP 5G-SRIT, 3GPP 5GRIT, and 5Gi (India/TSDSI). However, the 5Gi specification is unlikely to achieve widespread adoption outside of India[11]. Accordingly, the coverage in this article of the air interface standards is based on the 3GPP specifications. This article summarizes three key aspects of the air interface: antennas, physical layer, and channel coding.

Antennas

5G systems use *Multiple-Input/Multiple-Output* (MIMO) antenna systems extensively. Key features are base-station antennas consisting of large arrays of antennas and the use of *Beamforming*, and *Beam Management*.

In a MIMO scheme, the transmitter and receiver employ multiple antennas. The source data stream is divided into *n* substreams, one for each of the *n* transmitting antennas. The individual substreams are the input to the transmitting antennas (multiple input). At the receiving end, *m* antennas receive the transmissions from the *n* source antennas via a combination of line-of-sight transmission and multipath. The output signals from all of the *m* receiving antennas (multiple output) are combined. With a lot of complex math, the result is a much better receive signal than can be achieved with either a single antenna or multiple frequency channels. Note that the terms *input* and *output* refer to the input to the transmission channel and the output from the transmission channel, respectively.

MIMO systems are characterized by the number of antennas at each end of the wireless channel. Thus, an 8 × 4 MIMO system has eight antennas at one end of the channel and four at the other end. In configurations with a base station, the first number typically refers to the number of antennas at the base station. There are two types of MIMO transmission schemes:

1. *Spatial Diversity:* The same data is coded and transmitted through multiple antennas, effectively increasing the power in the channel proportionally to the number of transmitting antennas. This mechanism improves the *Signal-to-Noise Ratio* (SNR) for cell edge performance. Further, diverse multipath fading offers multiple "views" of the transmitted data at the receiver, thus increasing robustness. In a multipath scenario where each receiving antenna would experience a different interference environment, there is a high probability that if one antenna suffers a high level of fading, another antenna will have sufficient signal level.

2. *Spatial Multiplexing:* A source data stream is divided among the transmitting antennas. The gain in channel capacity is proportional to the available number of antennas at the transmitter or receiver, whichever is less. Spatial multiplexing can be used when transmitting conditions are favorable and for relatively short distances compared to spatial diversity. The receiver must do considerable signal processing to sort out the incoming substreams, all of which are transmitting in the same frequency channel, and to recover the individual data streams.

Beamforming is one of the essential technologies in developing advanced cellular antenna systems. Beamforming is a technique by which an array of antennas can be steered to transmit radio signals in a specific direction. Rather than simply broadcasting energy/signals in all directions, the antenna arrays that use beamforming determine the direction of interest and send/receive a stronger beam of signals in that specific direction.

In this technique, each antenna element is fed separately with the signal to be transmitted. The phase and amplitude of each signal are then added constructively and destructively in such a way that they concentrate the energy into a narrow beam or lobe. The various transmitted signals merge in the air by normal coherence of the electromagnetic waves, thereby forming a virtual beam in a predetermined direction. To understand how this procedure works, consider a signal that is fed to different antenna elements shifted in phase different amounts for each element. Now picture the transmitted energy from each element at an angle of 45°. At any point along that 45° line, the distance traveled by electromagnetic waves from different antenna elements is not equal. If the phase shifting is such that at 45° signals from all antenna elements arrive at the same phase, then the beam is strongest in that direction.

Beam Management refers to techniques and processes used to achieve the transmission and reception of data over relatively narrow beams. Beamforming and beam management are essential for using the *millimeter-wave* (mmWave) region over the 5G air interface. Narrow beams are needed to compensate for high path loss and blockage. With the use of narrow beams, and especially if the UE is mobile, beam management provides the means for both the base-station antenna and the onboard UE antenna to "lock on" to a beam that provides an optimal path from transmitter to receiver.

By adjusting the phase and amplitude parameters, a MIMO antenna can generate multiple beams, with each beam covering part of the cell area. For downlink transmission, the objective of beam management is to select a transmit beam to a UE so that the UE can receive the signal with the highest power and best SNR. For uplink transmission, the base station attempts to choose the receive beam for a UE with the best receive beamforming gain. Similarly, if the UE antenna system is capable of beamforming, the UE can use the beams to improve link quality.

The beam management procedure involves beamforming, beam sweeping, beam measurement, beam determination, and beam reporting, as shown in Figure 9.

This figure is taken from [12], which indicates the following elements in the context of downlink transmission:

- *Beam Sweeping:* The base-station antenna (that is, the 5G radio access network node gNB) transmits beams in a predetermined sequence for beam measurement at the UE side.
- *Beam Measurement:* The UE measures the characteristics of received beamformed signals.
- *Beam Determination:* The UE selects the optimal beam. In essence, the UE isolates the receive beam, which affords the best reception. The best results are obtained when the transmitting and receiving beam pair is optimal for the location of the UE at the time.
- *Beam Reporting:* The UE reports back to the gNB the information based on beam measurement.

Figure 9: Beam Management Procedures with Downlink Transmissions

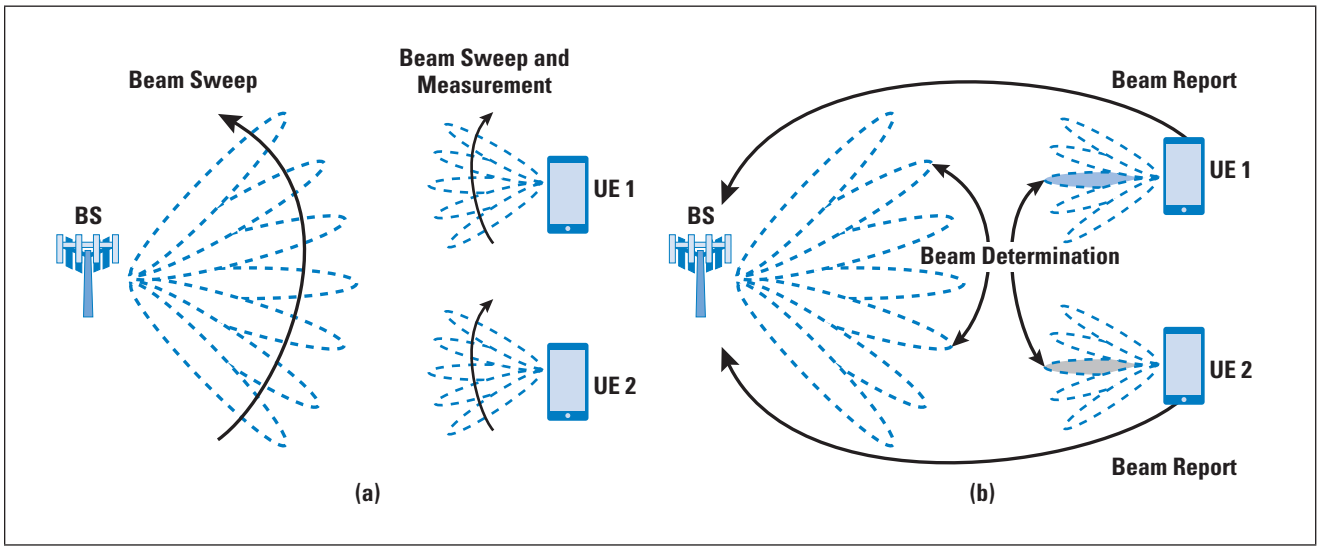

Beam management is an ongoing dynamic process that involves selecting an initial beam pair and then modifying the selection as transmit/receive conditions change.

The term *full-dimension MIMO* (FD-MIMO), or *3D-MIMO*, refers to a MIMO antenna system that is capable of varying the direction of a beam in both horizontal (azimuth) and vertical (elevation) dimensions. Thus, FD-MIMO can project a beam in any direction in three-dimensional space. This capability has advantages, especially in dense urban environments. The ability to adjust transmitted beams in the vertical dimension can improve the received signal power of terminals deep inside high-rise buildings and help overcome some of the building penetration loss. FD beamforming is also advantageous in indoor deployments in high-rise buildings, where a single base station may be able to optimize coverage over more than one floor. Such techniques directly increase spectral efficiency.

OFDMA and SC-FDMA

An important access of the air interface is the way multiplexing and multiple access is achieved over a physical transmission channel. Two techniques are employed for the 5G air interface: *Orthogonal Frequency Division Multiple Access* (OFDMA) and *Single-Carrier Frequency Division Multiple Access* (SC-FDMA). These two schemes use the following foundational techniques:

- *Frequency-Division Multiplexing* (FDM): A physical-layer technique in which multiple baseband signals are modulated on different frequency carrier waves and added together to create a composite signal. The effect of FDM is to divide a transmission bandwidth into multiple subchannels, each of which is dedicated to a particular baseband signal.

- *Frequency-Division Multiple Access* (FDMA): An access method at the data-link layer based on FDM principles, in which different frequency bands are allocated to different data streams, possibly from different users. The data-link layer in each station tells its physical layer to make a bandpass signal from the data passed to it. The signal must be created in the allocated band. There is no multiplexer at the physical layer. The signals created at each station are automatically bandpass-filtered. They are mixed when they are sent to the common channel. FDMA supports demand assignment, in which the assignment of frequency bands to users changes over time.

Orthogonal Frequency Division Multiplexing (OFDM), also called *Multicarrier Modulation*, is dedicated to a single data source. It uses multiple carrier signals at different frequencies, sending some of the bits on each channel. It differs from ordinary FDM in that the individual subcarriers are orthogonal to one another. In essence, signals are orthogonal to one another if the peaks of the power spectral density of each subcarrier occur at a point at which the power of other subcarriers is zero. A result of this property is that adjacent subcarriers can be packed closely together, making efficient use of the bandwidth.

OFDM has several advantages. First, frequency-selective fading affects only some subcarriers and not the whole signal. If the data stream is protected by a forward error-correcting code, this type of fading is easily handled. More importantly, OFDM overcomes *Intersymbol Interference* (ISI) in a multipath environment. ISI has a greater impact at higher bit rates because the distance between bits, or symbols, is smaller. With OFDM, the data rate is reduced by a factor of *N*, and this reduction increases the symbol time by a factor of *N*. This increase dramatically reduces the effect of ISI. As a result of these considerations, with the use of OFDM it may not be necessary to deploy equalizers, which are complex devices whose complexity increases with the number of symbols over which ISI is present.

Like OFDM, OFDMA employs multiple closely spaced subcarriers, but for OFDMA, the subcarriers are divided into groups of subcarriers. Each group is referred to as a *subchannel*. The subcarriers that form a subchannel need not be adjacent. In the downlink, a subchannel may be intended for different receivers. In the uplink, a transmitter may be assigned one or more subchannels.

Subchannelization defines subchannels that can be allocated to *Subscriber Stations* (SSs) depending on their channel conditions and data requirements. Using subchannelization within the same time slot, a *Base Station* (BS) can allocate more transmit power to user devices (SSs) with lower SNR, and less power to user devices with higher SNR. Subchannelization also enables the BS to allocate higher power to subchannels assigned to indoor SSs, resulting in better in-building coverage.

Subchannels are further grouped into *bursts*, which can be allocated to wireless users. Each burst allocation can be changed from frame to frame as well as within the modulation order. This capability allows the base station to dynamically adjust the bandwidth usage according to the current system requirements. Subchannelization in the uplink can save user-device transmit power because it can concentrate power on only certain subchannels allocated to it. This power-saving feature is particularly useful for battery-powered user devices, the likely case in mobile 4G and 5G.

SC-FDMA is a more recently developed multiple-access technique that is similar in structure and performance to OFDMA. One prominent advantage of SC-FDMA over OFDMA is the lower *Peak-to-Average Power Ratio* (PAPR) of the transmit waveform, which benefits the mobile user in terms of battery life and power efficiency. OFDM signals have a higher PAPR because, in the time domain, a multicarrier signal is the sum of many narrowband signals. At some time instances this sum is large, and at other times it is small, meaning the peak value of the signal is substantially larger than the average value. Thus, SC-FDMA is superior to OFDMA. However, it is restricted to uplink use because the increased time-domain processing of SC-FDMA would entail considerable burden on the base station.

Figure 10 provides an example of how the OFDM and SC-FDMA signals appear. As the figure illustrates, with OFDM, a source data stream is divided into N separate data streams, and these streams are modulated and transmitted in parallel on N separate subcarriers, each with bandwidth f_b. The source data stream has a data rate of R bps, and the data rate on each subcarrier is R/N bps. For SCFDMA, it appears from Figure 10 that the source data stream is modulated on a single carrier (hence the SC prefix to the name) of bandwidth $N \times f_b$ and transmitted at a data rate of R bps. The data is transmitted at a higher rate but over a wider bandwidth compared to the data rate on a single subcarrier of OFDM. However, because of the complex signal processing of SC-FDMA, the preceding description is not accurate. In effect, the source data stream is replicated N times, and each copy of the data stream is independently modulated and transmitted on a subcarrier, with a data rate on each subcarrier of R bps.

Compared with OFDM, SC-FDMA transmits at a much higher data rate on each subcarrier, but because the same data stream is on each subcarrier, it is still possible to reliably recover the original data stream at the receiver.

Figure 10: Example of OFDMA and SC-FDMA

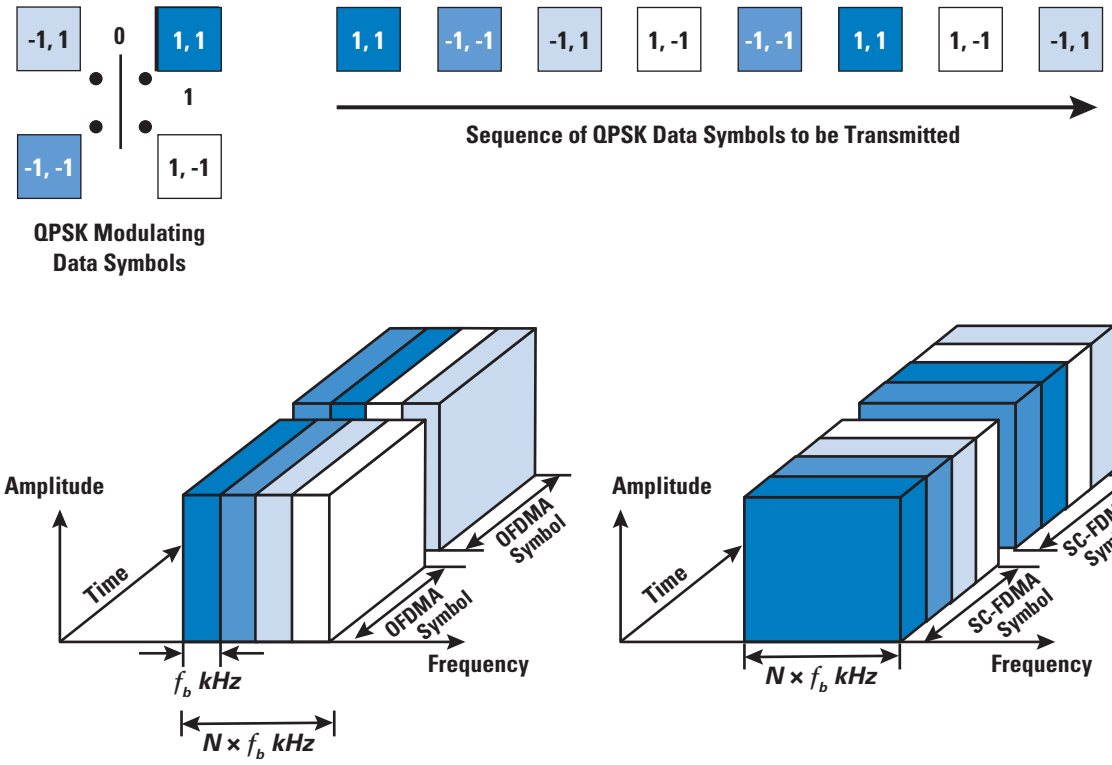

(a) OFDMA: Data symbols occupy f_b kHz for one OFDMA symbol period

(b) SC-FDMA: Data symbols occupy $N \times f_b$ kHz for 1/N SC-FDMA symbol period

Channel Coding

3GPP TS 38.212[13] specifies two *Forward Error Correction* (FEC) techniques for the air interface: *Low-Density Parity-Check Coding* and *Polar Coding*. TR 38.802[14] contains the results of a study into NR physical-layer aspects, and is useful for understanding the reasoning behind the concepts. An adequate overview of these two techniques is beyond the scope of this article, but a brief overview of the concepts is provided.

An (n, k) *parity-check code* encodes k data bits into an n-bit codeword. Typically, and without loss of generality, the convention used is that the leftmost k bits of the codeword reproduce the original k data bits, and the rightmost $(n - k)$ bits are the check bits. Such a code is defined by a set of $m = (n - k)$ simultaneous linear equations. If there are m linearly independent equations, there will be some set of k of the variables that can be arbitrarily specified such that one can solve for the other $(n - k)$ variables.

A parity-check code that produces n-bit codewords is the set of solutions to the following equations:

$$\begin{aligned} h_{11}c_1 \oplus h_{12}c_2 \oplus \ldots + h_{1n}c_n &= 0 \\ h_{21}c_1 \oplus h_{22}c_2 \oplus \ldots \oplus h_{2n}c_n &= 0 \\ &\bullet \\ &\bullet \\ &\bullet \\ h_{m1}c_1 \oplus h_{m2}c_2 \oplus \ldots \oplus h_{mn}c_n &= 0 \end{aligned}$$

…where the coefficients h_{ij} take on the binary values 0 or 1. The specific set of values of h_{ij} define a specific code.

The $m \times n$ matrix $H = [h_{ij}]$ is called the *Parity Check Matrix*. Each of the m rows of H corresponds to one of the individual equations. Each of the n columns of H corresponds to one bit of the codeword. If we represent the codeword by the row vector $c = [c_j]$, then the equation set can be represented as:

$$Hc^T = cH^T = 0$$

A *Low-Density Parity-Check* (LDPC) code is one in which H has a small density of 1s. That is, the elements of H are almost all equal to 0. Hence the designation *low density*. LDPC codes are enjoying increasing use in high-speed wireless specifications, including Wi-Fi, satellite, and cellular. LDPC codes exhibit performance in terms of bit error probability that is very close to the Shannon limit and can be efficiently implemented for high-speed use.

LDPC codes are suitable for larger blocks of data and are used for 5G data channels. Greater efficiency for small blocks of data is achievable with polar codes, and thus they are used for control channels. Polar codes involve a relatively complex mathematical transformation that involves splitting a communication channel into numerous synthetic bit channels, some of which have extremely low bit error rates and the remainder have high bit error rates, with the data bits being sent over the reliable bit channels. The mathematics behind this transformation is fairly complex and is not pursued here.

References and Further Reading

[0] William Stallings, "Introduction to 5G Part One: Standards, Specifications, and Usage Scenarios," *The Internet Protocol Journal*, Volume 26, No. 2, September 2023.

[1] 3GPP TS 23.501, "Technical Specification Group Services and System Aspects; System Architecture for the 5G System (5GS); Stage 2 (Release 16)," December 2020.

[2] 3GPP TS 23.502, "Technical Specification Group Services and System Aspects; Procedures for the 5G System (5GS); Stage 2 (Release 16)," December 2020.

[3] ITU-T, "Framework of software-defined networking," ITU-T Recommendation Y.3300, June 2014.

[4] ETSI, "Network Functions Virtualisation (NFV); Ecosystem; Report on SDN Usage in NFV Architectural Framework," ETSI GS NFV-EVE 005, December 2015.

[5] Xin Li, Mohammed Samaka, H. Anthony Chan, Deval Bhamare, Lav Gupta, Chengcheng Guo, and Raj Jain, "Network Slicing for 5G: Challenges and Opportunities," *IEEE Internet Computing*, September/October 2017.

[6] Next Generation Mobile Networks Alliance (NGMN Alliance), "5G End-to-End Architecture Framework," August 2019.

[7] GPP TS 38.300, "Technical Specification Group Radio Access Network; NR; NR and NG-RAN Overall Description; Stage 2 (Release 16)," January 2021.

[8] GSM Association, *The Mobile Economy*, published annually, **https://www.gsma.com/mobileeconomy/**

[9] 3GPP TS 38.401, "Technical Specification Group Radio Access Network; NG-RAN; Architecture Description (Release 16)," September 2020.

[10] ITU-R, "Detailed specifications of the terrestrial radio interfaces of International Mobile Telecommunications-2020," ITU-R M.2150, February 2021.

[11] Manu Kaushik, "India's own 5G standard could delay its 5G launch," *Business Today*, February 20, 2021.

[12] Guosen Yue, Lingjia Liu, Yongxing Zhou, and Jianzhong Zhang, "MIMO Technologies in 5G New Radio," *GetMobile: Mobile Computing and Communications*, March 2017.

[13] 3GPP TS 38.212, "Technical Specification Group Radio Access Network; NR; Multiplexing and channel coding (Release 16)," December 2020.

[14] 3GPP TR 38.802, "Technical Specification Group Radio Access Network; Study on New Radio Access Technology Physical Layer Aspects (Release 14)," September 2017.

[15] William Stallings, "Network Functions Virtualization," *The Internet Protocol Journal*, Volume 24, No. 2, July 2021.

[16] 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**

Why ATM Failed

by Craig Partridge, Colorado State University

In the late 1980s and early 1990s, *Asynchronous Transfer Mode* (ATM) was widely viewed as the new Internet architecture poised to take the place of the nascent Internet and to inaugurate a worldwide high-speed communications infrastructure. It didn't happen. Instead, after several years of uncertainty, the Internet swept ATM to the side and grew into the global infrastructure we know today.

Even today, 30+ years later, there are different views about how and why ATM "failed." This essay, while acknowledging ATM had to overcome some technical hurdles, argues that the central problem was a fast-moving window of opportunity that was squandered, largely because of poor standards leadership.

Origins of ATM

Jonathan Turner's forward-looking essay "New Directions in Communications (or Which Way to the Information Age?)," published in 1986, is widely viewed as the paper that launched ATM[1]. Turner examined the growing diversity of applications using data networks—in particular, advent of many-channel cable television. He looked at the rapidly diminishing error rates in transmission networks, thanks to the advent of fiber-optic cables.

Turner predicted that to meet future needs, our communications networks should be designed around high-performance parallel switches optimized for short packets of information sent over low-error links. Building on prior work at Bell Labs, Turner anticipated that to keep overhead (notably headers) in the short packets small, packets would contain small labels that associated the packets with established connections through the network (both point-to-point and point-to-multipoint).

That, simply, was the idea behind Asynchronous Transfer Mode: implement a futuristic network in which data was reliably transported at high speed in small, fixed-size packets called *cells* over connections.

Turner's central insights were right. Today's data communications networks are built around highly reliable transmission networks and make heavy use of high-performance parallel switches. Furthermore, those high-performance parallel switches internally move data in ways akin to Turner's proposed small packets.

When Turner made his predictions, the Internet was entering one of most difficult stages in its explosive growth. In late 1986, the Internet became plagued by congestion and routing collapses and struggled with inadequate network-management capabilities.

The Internet standards process was not up to the task of handling these concurrent challenges and, as a result, it was necessary to create the *Internet Engineering Task Force* (IETF) to coordinate the substantial efforts to convert research into standardizations. It would take until 1989 for the critical issues to get resolved and their solutions standardized.

Thus, in 1986, future-focused people in data networking saw a lot of wisdom in Turner's paper and a level of worry when they looked at the state of the Internet. Planning for a better data-communications future based on Turner's ideas was appealing.

Nonetheless, ATM, the realization of Turner's vision, failed for three reasons. First, ATM suffered from exceptionally poor standards leadership. Second, and due in part to the poor standards leadership, ATM missed the window(s) of opportunity to capture the local-area network market. Finally, ATM had difficulties matching the needs of the emerging wide-area Internet market.

The ATM Choices in 1986

Before delving into where ATM failed, it is useful to look at the futures one could envision for ATM in 1986 after reading Turner's paper. Turner suggested at least three choices, and the choices were not exclusive:

- ATM as the future universal data-networking protocol. This choice was the most intellectually popular one and, furthermore, the one the telephony industry wanted to see. In this plan, ATM, with its high performance, would sweep away the various networking alternatives such as TCP/IP, *Open Systems Interconnection* (OSI), Apple's *AppleTalk*, IBM's *Systems Network Architecture* (SNA), and Novell *Netware*. ATM would be end-to-end, from the wall jack in your home or office to the data center or business. Furthermore, ATM could also support voice and video—in ways the Internet could not yet—thus positioning the telephone industry to retain their existing voice business and take customers from the cable TV industry.
- ATM as a *Local-Area Network* (LAN) technology[2]. In 1986, the state of local-area networks was poor. The major technology was original Ethernet, which required heavy coax cables. Interconnecting Ethernets was tedious and required hand configuration to avoid routing loops. It was just in 1986 that Digital Equipment Corporation introduced the *Spanning Tree Protocol*, which prevented loops and made it much easier to connect Ethernet segments. Thus, in 1986, LANs looked clunky, and ATM looked like a way to make it easier to build large corporate networks. This approach was particularly appealing to Silicon Valley startups, because it was relatively simple to build a 4- or 8-port ATM switch, and they could envision that if their small switches sold well, they would be positioned to move into the wide-area ATM market when that market matured.

- Finally, you could envision ATM as just-another-link layer over which you ran TCP/IP. At its simplest, the idea was that ATM would be the next-generation *Wide-Area Networking* (WAN) protocol and TCP/IP would run on top of it. This idea was popular primarily within the burgeoning TCP/IP community, but it was also acceptable in Silicon Valley, where the ATM-as-a-LAN product vendors were happy to have customers for their products, even if those customers used TCP/IP. It was, of course, anathema to the telephony community, which sought to use ATM to take control of the growing data-communications market.

Exceptionally Poor Standards Leadership

Turner's vision of the future was particularly appealing to the international telephony community, which had missed the early stages of the data-communications revolution and the cable-television revolution. Telephone companies, many of which were quasi-government owned, were (correctly) concerned that their business model centered on voice communication was going to be destroyed and their business would shrink to managing the fiber and cables over which other companies would make money selling data services.

So, the telephone companies tasked a committee of their standards organization—the *International Telecommunication Union* (ITU)—to make Turner's vision a reality. The *Consultative Committee for International Telegraphy and Telephony* (CCITT) began creating ATM. (In 1992, CCITT became the *International Telecommunication Union Telecommunication Standardization Sector* [ITU-T]).

CCITT was utterly unsuited to create data-communications standards. It was a standards group that drew its expertise largely from telephony laboratories. A senior US data-communications researcher from *Bell Communications Research* ["Bellcore"](which had spun off from Bell Labs) attended one of the early ATM standards meetings, assessed the data-communications expertise in the room, and promptly announced to his friends that ATM was an acronym for "Another Telephony Mistake."

CCITT launched straight into standards making without doing an architectural review. Turner had assumed a world of high-speed, low-error networks, and those assumptions had architectural implications. For example, in 1990, Julio Escobar and I did a detailed study of how to do error detection and recovery for ATM in fiber-optic networks. When Julio took our results to a CCITT meeting to figure out what we had missed, he was stunned to be welcomed as the first person who had done an analysis.

CCITT assumed that the world in which the telephone, the cable TV, and the computer each had its own distinctive wall plug would persist in ATM. This assumption was at odds with Turner's vision. In his paper, Turner had pointed out that application-specific networks were a mistake.

It was widely understood in the data-communications community that all data was *bits* and, for the most part, those bits did not care about the format or the wire they were transmitted over. But CCITT's member telephone companies wanted to charge different tariffs for different services, so having a uniform protocol for all data was undesirable. Rather the notion was they would charge distinct tariffs for voice, video, and data, and enforce this differentiation by formatting the different types of data in distinct formats in ATM. Accordingly, CCITT went ahead and defined distinct cell formats, called *ATM Adaptation Layers* (AALs), for voice, video, and data. Each AAL was mapped into a standard 53-byte cell format.

The development of the AAL for data was a particular disaster. The standards group initially planned for two data AALs, AAL 3 for digital video and AAL 4 for computer data. They realized the two could be combined, and they created AAL 3/4, which was swiftly standardized, but was so badly designed that data-communications-savvy members of the standards committee consulted with members of the IETF to propose a new AAL for data, which became AAL 5[3].

Finally, that 53-byte cell size is worth attention. It came about because of a disagreement about data rates. Telephone companies with less-developed networks were planning to offer ATM over 1.5- or 2-Mbps links (US T1, European E0), and they wanted to support voice calls over ATM. To avoid jitter, that support dictated a small cell size. In contrast, companies in the US expected to offer ATM on 155-Mbps (OC-3) or greater high-speed links and wanted larger cell sizes to reduce the cost of fragmenting data into cells. Competing proposals of 16- and 128-byte cells moved to 32 and 64 bytes, respectively, and the compromise was 48 bytes plus a 5-byte header. The result was to meet neither party's goals.

By late 1991, the failings of the ATM standards process were so severe that the emerging ATM vendor community had to step in. The vendors announced an industry-driven standards group, *The ATM Forum*. The Forum effectively took over the ATM standards process.

A Window of Opportunity in the Local Area

In the grand telephony vision, ATM was designed to be a new networking technology, delivered by the telephone company to your office or home. To achieve that dream, ATM needed to win the competition for the home and office network. In 1986, when ATM was first envisioned, its adoption looked eminently possible, because there was no home network (people dialed into their office computer [not the network, but a specific computer] using a modem) and the primary office network technology was the cumbersome original Ethernet.

But change was also coming. The first thin-wire Ethernet standard, 802.3e, was standardized in 1986. The standard for 10BASE-T, which worked over twisted pair, would come in 1990.

With the advent in 1989 of the first multi-port Ethernet switches (from Kalpana), setting up an office Ethernet became a matter of running 4-wire cables to a closet that held an Ethernet switch. Physically that was the same service ATM was planning to offer, and thanks to Ethernet bridging standards, it was easier to install and operate Ethernet than the nascent ATM LANs.

Consumers could see that more Internet-compatible technology was coming soon. Plans for *100BASE-T* (100-Mbps Ethernet) were soon well-known. It came out in 1995. Wireless networking appeared in 1990 (*WaveLAN*), and standardization efforts leading to *Wi-Fi*, which was intentionally made easy to integrate with wired Ethernet, soon followed.

Concurrently, the Internet was booming. Its major technical problems having been resolved, its user population grew 16X between 1990 and 1995. The World Wide Web appeared in 1991. The Internet was built around Ethernet at the edges and long-haul leased lines in the core. If you had joined the Internet revolution, ATM meant changing your working installed technology. If you could convince yourself that the Ethernet growth curve was good enough, then you didn't need ATM LANs.

Thus, by late 1990, ATM had serious market competition at the edge. It was clear the competition was going to increase. There was a narrow opportunity (perhaps already lost) to capture market share and make ATM the networking service at the network edge. The realization the market window was closing helps explain the vendors' desire to fix the ATM standards process and create The ATM Forum in 1991.

In retrospect, ATM never managed to grab that market window, and the advent of 100-Mbps Ethernet slammed the window shut. Corporate customers had bought a handful of ATM switches to see if they might work if Ethernet and Wi-Fi did not evolve fast enough. But Ethernet and Wi-Fi did evolve fast enough. Furthermore, experience with ATM LANs did not make a compelling case for change.

A Last Chance in the Wide Area

While ATM was rapidly losing credibility for the office and home LAN market at the start of the 1990s, it still had a viable potential role as the technology for wide-area networks. The Internet relied on telecommunications companies for its long-haul links. It was entirely possible that those long-haul links could be running ATM.

At the start of the 1990s, no one knew how to build a multigigabit Internet router. The only working devices that could move data at line speed between multiple gigabit links were ATM switches. Admittedly the switches were prototypes and vendors were waiting for ATM standards to finalize the products, but realized versions of those devices existed, whereas a multigigabit router did not.

So, in many ways, a version of Turner's vision was still alive. The difference from 1986 was that rather than seeing the whole network as running something like ATM, the vision was that there would be islands of Ethernets running IP interconnected via ATM (with no routers in the network center because the routers were too slow).

As in the LAN market, emerging technologies and market forces meant ATM needed to move swiftly to grab this opportunity. In defense of the folks working on ATM for the wide area, unlike the LAN market (where the competitive situation was obvious to anyone who wished to look), the competition for the wide area was not widely recognized. It looked as if there was plenty of time to make this vision happen. As a result, just as in the LAN market, the wide-area ATM effort did not move fast enough.

Critical to this portion of the ATM story is the rise of businesses that made their money installing and selling long-haul runs of fiber-optic links. A British company, *Cable and Wireless*, jumped into this business when AT&T was broken up in 1984, and other companies, including what would become Level 3, followed. These companies started by renting fiber-optic links to the telephone carriers, but in the early 1990s they realized there was a possible market selling links to *Internet Service Providers* (ISPs).

The links ran a point-to-point protocol using the *Synchronous Optical NETwork* (SONET) or *Synchronous Digital Hierarchy* (SDH) protocols (which differed in minor details). SONET/SDH deliver blocks of bytes at a range of speeds from 155 Mbps up to many gigabits.

By mid-1993, there was a draft of a standard for running IP (and other protocols) over ATM. This standard was recognized to be at least a year later than the market needed it. In this case, arguments among vendors had caused the IETF to take too long. Nonetheless, when the standard was issued, the telephone industry could have sold high-speed (155 Mbps and faster) static ATM circuits running to ISPs. There was early adopter interest in such a product. But, for reasons unclear, the telephone industry did not offer a product.

Enterprising technologists realized there was an opening to develop a competing product to meet ISPs' need to be able to transmit IP datagrams over SONET links. In mid-1994, a draft of a standard for the *Point-to-Point Protocol* (PPP) operating over SDH/SONET appeared. PPP over SONET/SDH was an appealing product because ISPs could simply lease a fiber between two points of presence and run PPP. That approach allowed the ISPs to bypass the telephone companies by renting fiber directly and running PPP. An added benefit was that PPP made better use of the link (less overhead) than ATM and was simpler to manage.

At this point, the ATM dream was nearing its end. The only remaining hope was that the TCP/IP protocols would fail to scale cleanly to gigabit speeds.

By 1994, the only challenge remaining for TCP/IP was the development of multi-gigabit (10+ Gbps and faster) Internet routers. But by early 1995, multiple router vendors were indicating to their customers that, while challenges remained, the customers should expect multi-gigabit IP routers to appear.

Conclusion

In 1996, ten years after Turner's paper, ATM as a forward-looking networking protocol was effectively dead. The IETF's IP over ATM working group held its last meeting in March 1996.[4,5,6] Perhaps more vividly illustrating the situation, a startup named Ipsilon was marketing a product that sought to repurpose ATM switches as IP routers.

ATM had its origins in Jon Turner's clear vision of the future, and it had a 2- to 5-year head start on the protocols it would compete with. This essay suggests that the ATM community squandered that lead. A valid alternative explanation is that ATM was too complex. Creating a high-speed, circuit-switched, data-networking protocol was a difficult problem in the 1980s, and ATM struggled with challenges such as address resolution and congestion control. But those challenges were ultimately solved. I suggest they would have been solved earlier had the ATM community felt more urgency, and thus this essay focuses on missed chances rather than a technical reason for the ATM failure.

References and Further Reading

[1] Jonathan Turner, "New Directions in Communications (or Which Way to the Information Age?)," in *IEEE Communications Magazine*, Volume 24, No. 10, pp. 8–15, October 1986.

[2] J. Bryan Lyles and Daniel C. Swinehart, "The emerging gigabit environment and the role of local ATM," in *IEEE Communications Magazine*, Volume 30, No. 4, pp. 52–58, April 1992.

[3] Daniel H. Greene and J. Bryan Lyles, "Reliability of Adaptation Layers," in *Proceedings of the IFIP WG6.1/WG6.4 Third International Workshop on Protocols for High-Speed Networks*, Stockholm, Sweden, May 13–15, 1992. IFIP Transactions C-9, North-Holland 1993, ISBN 0-444-89925-1.

[4] Mark Laubach, "Classical IP and ARP over ATM," RFC 1577, January 1994.

[5] Mark Laubach and Drew Perkins, "IP over ATM Working Group's Recommendations for the ATM Forum's Multiprotocol BOF Version 1," RFC 1754, January 1995.

[6] Robert G. Cole, David H. Shur, and Curtis Villamizar, "IP over ATM: A Framework Document," RFC 1932, April 1996.

[7] George Clapp and Mike Zeug, "Components of OSI: Asynchronous Transfer Mode (ATM) and ATM Adaptation Layers," *ConneXions—The Interoperability Report*, Volume 6, No. 4, April 1992.
https://archive.org/details/ConneXions.06.04

[8] Tom Lyon, "Simple and Efficient Adaptation Layer (SEAL)," ANSI Standards Project T1S1.5 AAL, August 1991.

[9] John T. Lewis, Raymond Russell, Fergal Toomey, Brian McGurk, Simon Crosby, and Ian Leslie, "Practical connection admission control for ATM networks based on on-line measurements," *Computer Communications*, Volume 21, Issue 17, pp. 1585–1596, November 25, 1998.

[10] Alexander G. (Sandy) Fraser, "Towards a Universal Data Transport System," in *IEEE Journal on Selected Areas in Communications*, Volume 1, No. 5, pp. 803–816, November 1983.

CRAIG PARTRIDGE is a professor at Colorado State University (CSU). Prior to coming to CSU, he was Chief Scientist at BBN Technologies. A member of the *Internet Hall of Fame*, Craig was one of the founding members of the *Internet Engineering Steering Group* (IESG), and he designed MX resource records. Craig can be reached at: **craig.partridge@colostate.edu**

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.

Lessons Learned from 20 Years of Cellular and Wi-Fi Integration

by Mark Grayson, Cisco Systems

By any measure, cellular and Wi-Fi wireless technologies can be viewed as having fundamentally transformed the way users consume Internet services. The number of cellular users worldwide now exceeds 7 billion[1], and the number of Wi-Fi devices shipped now exceeds 39 billion[2]. These numbers reflect the complementary nature of the wireless systems, with the exclusively licensed cellular systems enabling cellular operators to provide Internet access over wide geographic coverage and the shared unlicensed Wi-Fi systems enabling businesses and individuals to provide targeted local-area coverage.

Whereas the default approach is for these two wireless systems to simply co-exist while supporting their respective value propositions, there have been various attempts over the last 20 years to integrate the two technologies into a single "converged" architecture. Over those 20 years, and the different "Gs," a myriad of architectural approaches has been specified by the *3rd Generation Partnership Project* (3GPP) to integrate Wi-Fi and cellular architectures.

This article looks at some of the key takeaways from the last 20 years of attempting to converge cellular and Wi-Fi systems.

3GPP Specifications

Efforts started in December 2003, when 3GPP approved its first work item for "UMTS-WLAN Interworking"[3]. The justification behind the work item highlighted the complementary nature of cellular and Wi-Fi deployments:

> *"WLAN technology can complement 3GPP based networks in deployment environments with high user density and demand for higher data rates. However, in order to provide flexible use of both technologies in these environments and to provide mobility of services between the two technologies it is sensible that some degree of interworking exists between the two technologies/systems."*

Fast forward 20 years and there now have been over a dozen different approaches specified that look to "converge" cellular. These are listed in Table 1, with some of the architectures illustrated in Figure 1.

Table 1: A range of 3GPP approaches for converging 3GPP and Wi-Fi systems.

2G-Based	Generic Access Network
3G-Based	Interworking WLAN
4G-Based	Access Network Discovery and Selection Function, evolved Packet Data Gateway, Dual Stack Mobile IPv6, IP Flow Mobility and Seamless WLAN offload, Proxy Mobile IPv6 Trusted WLAN, S2a GPRS Tunnelling Protocol Trusted WLAN, LTE WLAN Aggregation, LTE/WLAN Radio Level Integration with IPsec Tunnel, Network-Based IP Flow Mobility
5G-Based	Non-3GPP Interworking Function, 5G Access Traffic Steering Switching and Splitting

Whereas the majority of these architectures have failed to see wide-scale adoption, the one solution that has seen significant deployment is the evolved *Packet Data Gateway* (ePDG) used to support *Wi-Fi Calling*, illustrated in Figure 1.

Although standardized in 2008[4], it wasn't until 2014 with the launch of Apple iOS 8 and its native Wi-Fi Calling capability that the functionality became widely adopted, allowing transparent access to *IP Multimedia Subsystem* (IMS)-based rich media communications over both *Long-Term Evolution* (LTE) and Wi-Fi access networks. This type of deployment leverages enhanced *User Equipment* (UE) functionality to use an *IP Security* (IPSec) tunnel between the UE and ePDG to support the IMS-based services. Figure 1 shows the architecture of 3GPP.

Figure 1: 3GPP's ePDG Architecture

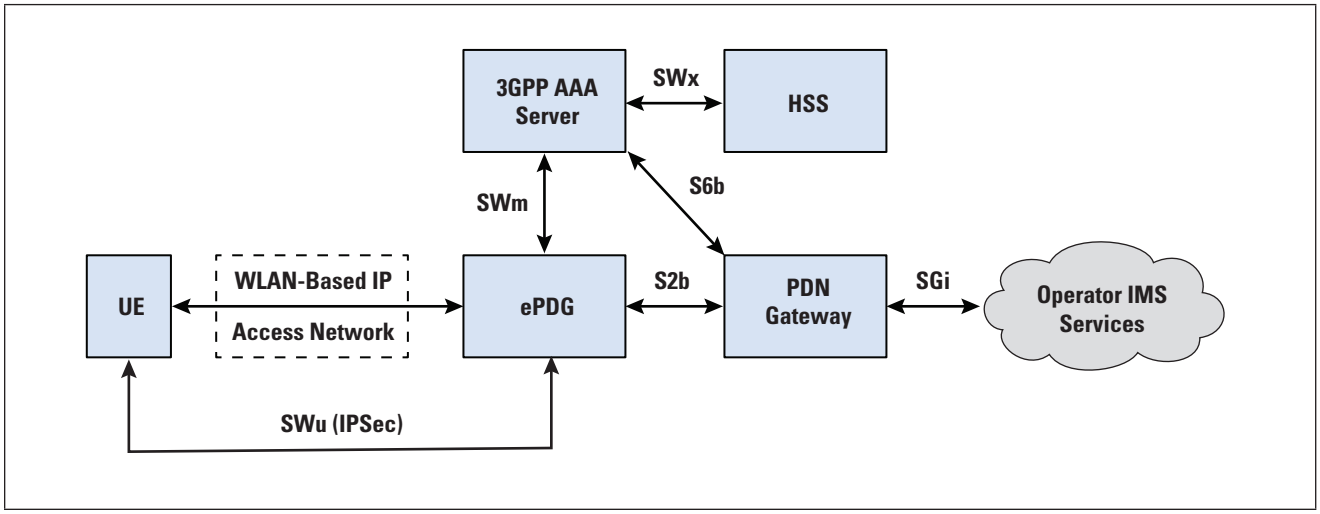

GSMA IR.67[5] and 3GPP TS23.003[6] have defined a standard realm that mobile operators may use in their Wi-Fi Calling deployments to enable their ePDGs to be discoverable over the public Internet. The *Fully Qualified Domain Name* (FQDN) is of the form:

```
epdg.epc.mnc<MNC>.mcc<MCC>.pub.3gppnetwork.org
```

…where **<MCC>** represents an E.212 *Mobile Country Code* and **<MNC>** represents the E.212 *Mobile Network Code* allocated to the mobile operator.

Wi-Fi Calling Adoption

In June 2023, the list of allocated MCC and MNC values published at **http://mcc-mnc.com/** was used to determine whether the operator that had been allocated a particular MCC and MNC had configured a *Domain Name System* (DNS) entry to enable its ePDG to be discovered. Table 2 shows the results, which indicate that over 100 countries have deployed Wi-Fi Calling where the ePDG is discoverable using the standard FQDNs defined by 3GPP.

Table 2: Countries where at least one operator has configured a standard DNS entry for ePDG discovery (Source: Cisco Systems)

Albania	Colombia	Iceland	Morocco	Saint Lucia
Anguilla	Croatia	India	Myanmar	Saint Vincent and the Grenadines
Antigua and Barbuda	Cyprus	Indonesia	Namibia	Saudi Arabia
Argentina	Czech Republic	International Networks	Nepal	Singapore
Armenia	Denmark	Ireland	Netherlands	Slovakia
Australia	Dominica	Israel	New Zealand	Slovenia
Austria	Dominican Republic	Italy	Norway	South Africa
Bahamas	Ecuador	Jamaica	Oman	Spain
Bahrain	Egypt	Japan	Pakistan	Sri Lanka
Bangladesh	Estonia	Jordan	Panama	Sudan
Barbados	Faroe Islands	Kazakhstan	Paraguay	Sweden
Belarus	Finland	Kuwait	Peru	Switzerland
Belgium	France	Latvia	Philippines	Taiwan
Brazil	Germany	Liechtenstein	Poland	Thailand
British Virgin Islands	Ghana	Lithuania	Portugal	Türkiye
Brunei	Greece	Luxembourg	Puerto Rico	Turks and Caicos Islands
Bulgaria	Grenada	Malaysia	Qatar	Ukraine
Cambodia	Guadeloupe and Martinique and French Guiana	Maldives	Reunion	United Arab Emirates
Canada	Guam	Monaco	Romania	United Kingdom
Cayman Islands	Hong Kong	Montenegro	Russia	United States of America
Chile	Hungary	Montserrat	Saint Kitts and Nevis	Vietnam

There is clearly a disparity in adoption of the different 3GPP approaches for converging 3GPP and Wi-Fi systems. For instance, in contrast to the over 100 countries that have launched ePDG-based integration, the *Global Mobile Suppliers Association* (**www.gsacom.com**) reports that only a single operator has invested in LTE WLAN Aggregation.

IMS-based Wi-Fi Calling Observations

Compared to the alternative "trusted" solutions defined by 3GPP for integrating Wi-Fi, the ePDG-based integration can leverage any suitable Wi-Fi network. The result of this leverage may be one of the key reasons that has led to its rapid adoption.

When looking at the Wi-Fi market as a whole, Dell'Oro reports that around 6% of all Wi-Fi equipment revenue is associated with the Service Provider segment[7]. Only Manufacturing and Logistics segments have lower overall market share, with the Wi-Fi markets for K-12 Education, Higher Education, Finance, Healthcare, Government, Hospitality, and Retail all exceeding the Wi-Fi Service Provider market.

The first lesson learned is to avoid restricting your target market. By enabling all segments deploying Wi-Fi to benefit from ePDG-based integration, the Wi-Fi Calling approach offers the broadest market reach.

The next key observation is that the majority of smartphone data is being sent over connections that use Wi-Fi rather than mobile networks (2G, 3G, 4G, or 5G). The latest data from UK regulator *Ofcom* indicates that nearly three-questers of all smartphone data is sent over Wi-Fi rather than mobile[8]. Increasingly these Wi-Fi systems are being dimensioned to deliver gigabit-based services over the fixed network. When comparing data from Ofcom's latest *Communications Market Report*[9], the average volume for fixed broadband, where Wi-Fi dominates, is 453 GB a month, which is 75 times the 6 GB a month for the average cellular subscription.

These figures mean that the smartphone traffic transported over Wi-Fi equates to around 1% of the total fixed broadband traffic that can easily be accommodated. Equally important is the focus of ePDG-based integration on delivering seamless connectivity for IMS-based services, enabling users to receive mobile calls when out of cell tower coverage. With Ofcom reporting that the average UK cellphone user calls for 200 minutes of use a month, and a conservative 128 Kbps for the IMS call over Wi-Fi, the impact of Wi-Fi Calling on the cellphone network can be estimated:

Total 200 minutes (cellular and Wi-Fi) * 75% = 150 minutes over Wi-Fi
150 × 60 = 9000 seconds
9000 × 128 Kb = 1.1Gb = 144 MB

Importantly, the 144 MB/month of voice-over-Wi-Fi traffic corresponds to a 2.5% traffic increase compared with the average 6 GB/month used by a cellular subscriber.

This information can be contrasted with other integration approaches that focus on "trusted integration," where all traffic sent over Wi-Fi is integrated into the cellular provider's gateway. With 75% of traffic being carried over Wi-Fi, these approaches may result in a 300% increase in traffic across the cellular network. Whether the cellular operators can derive sufficient value from the 300% increase in traffic to cover the additional costs in supporting such is an open issue. However, the increasing adoption of encrypted flows over the Internet has already impacted an operator's ability to derive value from observing data sent over cellular networks.

The second lesson learned is to avoid thinking of Wi-Fi and cellular as symmetrical services. Wi-Fi is already being dimensioned to support 75 times the traffic load of cellular, and the majority of smartphone traffic continues to be carried on Wi-Fi. Hence, there appears to be advantages to focus integration efforts on systems that avoid transporting the bulk of Wi-Fi data over cellphone networks, such as enabled by IMS-focused ePDG-based integration.

Integrating Native IP-based Services

In the 20 years since 3GPP embarked on the journey to converge 3GPP and Wi-Fi, there has been a significant transition in how Internet services are consumed. Early attempts at convergence were hampered by the binding of sockets to physical interfaces, with applications often stalling as devices made the switch from cellular to Wi-Fi. Hence, initial architectures looked to mask transitions from client-side applications, including the use of Mobile IP client functionality that bound sockets to logical instead of physical interfaces.

In 2023, the Internet is continuing to transition. Not only is over 90% of Internet traffic encrypted, in certain regions of the world we observe that nearly 50% of the traffic has transitioned from regular TCP to HTTP3 transported over *User Datagram Protocol* (UDP)-based QUIC[10]. Critically, instead of having to mask different paths, the QUIC transport protocol supports native connection migration. Existing connections continue to operate as devices change their endpoint IP addresses when they switch between different networks.

Since 2022, hyperscaler offerings have included native support for HTTP3 and connection migration capability, and the device ecosystem has similarly enabled application developers to benefit from it[11].

The third lesson learned is to avoid thinking of situations where multiple accesses and multiple paths are available to devices as peculiar. Convergence solutions shouldn't be a "bolt-on" to address specific corner cases. Instead, accept that the Internet has already started its transition to natively support such scenarios.

The Complexities of Path-Selection Policy

Path-selection policy in a heterogeneous environment is a complex issue. Instead of the network-controlled handover approach used in homogeneous cellular networks, the characteristics of Wi-Fi and cellular connections may vary dramatically in terms of costs, quality, and, for moving users, coverage persistence. However, some of the convergence architectures look to expand service providers' cellular network-controlled approach to accommodate Wi-Fi, enabling service providers to define rules that include packet-flow descriptors, access-selection criteria, as well as how to control the steering of flows between Wi-Fi and cellular.

But there is now increasing acceptance that the network provider is but one stakeholder in the complex decision process that is path selection. Identity providers may have preferred relationships that lead them to prioritize the usage of specific paths.

Users are important stakeholders in path selection, including whether the path corresponds to an unmetered private connection or a metered network that may lead to additional charges. Operating System and device vendors can base their preference on near real-time visibility of access networks and associated metrics, and battery levels can be used to drive preference for paths that consume lower energy. Application providers know the metrics that result in the best application experience, whether that is lowest latency for interactive applications or highest throughput for applications that consume significant amounts of data. Applications know in advance the likely duration of application flows and whether it is worth migrating already established flows or waiting for the establishment of new flows over a newly preferred path.

The device ecosystem is looking to meet the needs of their application providers by delivering frameworks that enable applications to configure how multiple paths should be employed, enabling applications developers to easily benefit from the HTTP3 connection migration capability.

The final lesson learned is that a single command-and-control approach to path-selection policy cannot accommodate all stakeholder requirements. And we should recognize that value is continuing to migrate towards the application; application loyalty is the new brand loyalty. So, the goal should be about how to best deliver those application experiences, and what hints and instrumentation can be exchanged between stakeholders to enable better decisions to be made.

Summary

The last two decades have seen significant investment and innovation in the development of cellular and Wi-Fi integration for the delivery of enhanced mobile services. This article has looked at the some of the key takeaways from the journey and the lessons learned along the way. In summary: embracing an approach that facilitates integration with the 94% of non-service provider Wi-Fi deployments and leverages the native connection migration support provided by HTTP3, while ensuring that application stakeholders can exchange hints to enable better decisions, is the best route for delivering enhanced services across combined Wi-Fi and cellular networks.

References and Further Reading

[1] Petroc Taylor, "Forecast number of mobile users worldwide from 2020 to 2025," *Statista*, January 18, 2023.
https://www.statista.com/statistics/218984/number-of-global-mobile-users-since-2010/

[2] "Value of Wi-Fi," *Wi-Fi Alliance*,
https://www.wi-fi.org/discover-wi-fi/value-of-wi-fi

[3] "3GPP system - WLAN- Interworking," 3GPP Technical Specification Group Services and System Aspects, Meeting #22, Maui, Hawaii, USA, 15–18 December 2003,
https://www.3gpp.org/ftp/tsg_sa/TSG_SA/TSGS_22/Docs/PDF/SP-030712.pdf

[4] "3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Architecture enhancements for non-3GPP accesses (Release 12)," March 2013, **https://www.3gpp.org/ftp/Specs/archive/23_series/23.402/23402-c00.zip**

[5] "DNS Guidelines for Service Providers and GRX and IPX Providers, Version 21.0," *GSM Association*, 25 November 2022, **https://www.gsma.com/newsroom/wp-content/uploads//IR.67-v21.0.pdf**

[6] "3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Numbering, addressing and identification; (Release 18)," March 2023. **https://www.3gpp.org/ftp/Specs/archive/23_series/23.003/23003-i10.zip**

[7] "Market Research Reports on Wireless LAN," Dell'Oro Group, **https://www.delloro.com/market-research/enterprise-network-infrastructure/wireless-lan/**

[8] "Mobile Matters," *Office of Communications* (Ofcom), **https://www.ofcom.org.uk/research-and-data/telecoms-research/mobile-smartphones/mobile-matters**

[9] "Multi-sector Research," *Office of Communications* (Ofcom), **https://www.ofcom.org.uk/research-and-data/multi-sector-research/**

[10] Andreas Enotiadis and Bart Van de Velde, "The New, Encrypted Protocol Stack Taking over the Internet and How to Deal with It," *Cisco Live*, Las Vegas, June 2023. **https://www.ciscolive.com/on-demand/on-demand-details.html?#/session/1686177808340001V6Dc**

[11] Channy Yun, "New – HTTP/3 Support for Amazon CloudFront," *AWS Blog*, August 15, 2022. **https://aws.amazon.com/blogs/aws/new-http-3-support-for-amazon-cloudfront/**

[12] Geoff Huston, "A Quick Look at QUIC," *The Internet Protocol Journal*, Volume 22, No. 1, March 2019.

[13] Geoff Huston, "Comparing TCP and QUIC," *The Internet Protocol Journal*, Volume 25, No. 3, December 2022.

MARK GRAYSON holds an B.Eng. in Electronic and Communication Engineering from the University of Birmingham (England), and a Ph.D. in Wireless Communications from University of Hull. Since 2000, he has worked for Cisco Systems on a variety of wireless-related projects, including early EAP-SIM implementations, 3GPP standardization of Inter-working WLAN and Cisco's Service Provider Wi-Fi Solutions. He has served on the board of the *Small Cell Forum* where he was responsible for the small cell virtualization efforts that led to the specification of the nFAPI-based split 6. He now serves as co-chair of *O-RAN Alliance's Fronthaul Working Group*, and is rapporteur for the deliverables dealing with the YANG-based management of the O-RAN Radio Units. He is also the chair of the *Wireless Broadband Alliance's OpenRoaming Wi-Fi Federation* aimed at lowering barriers to roaming onto private networks and is a Fellow of the *Institute of Engineering and Technology* (IET). He can be reached at: **mg@cisco.com**

A New and Simplified Way to Request Nonpublic gTLD Registration Data

By Adiel Akplogan, ICANN

The *Internet Corporation for Assigned Names and Numbers* (ICANN) has recently launched the *Registration Data Request Service* (RDRS). This new service handles requests for access to nonpublic registration data related to *generic Top-Level Domains* (gTLDs). The RDRS is a free and global service that can be an important resource for ICANN-accredited registrars and those who have a legitimate interest in nonpublic data like law enforcement, intellectual property professionals, cybersecurity professionals, consumer-protection advocates, and government officials. The service introduces a more consistent and standardized format for handling these unique requests.

Because of personal data protection laws, many ICANN-accredited registrars are now required to redact personal data from public records such as WHOIS[1] lookups. With no one way to request or access such data, it can be difficult for interested parties to get the information they need. The RDRS will help ease this problem by providing a simple and standardized process to make these types of requests with benefits for both registrars and requestors.

The RDRS is not only an important tool for the Internet community at large but for the ICANN Board as well. The service was implemented at the direction of the ICANN Board to gather relevant usage data to help inform policy decisions related to a *System for Standardized Access/Disclosure*[2]. The more registrars and requestors that use the RDRS, the more accurate and valuable the data collected will be toward making that decision. ICANN-accredited registrars are encouraged to opt-in to the service. More information is available at the end of this article.

What Is the RDRS?

The RDRS is a free, global, one-stop-shop ticketing system that handles nonpublic gTLD registration data requests. The RDRS connects requestors of nonpublic data with the relevant ICANN-accredited registrars for gTLD domain names that participate in the service. The service streamlines and standardizes the process for submitting and receiving requests through a single platform. It is important to note that the RDRS does not guarantee access to requested registration data. All communication and data disclosure between the registrars and requestors takes place outside of the system.

Who Can Use the RDRS?

The service is intended for use by ICANN-accredited registrars and individuals and entities with a legitimate interest for access to nonpublic gTLD registration data.

Requestors include but are not limited to: law enforcement, intellectual property professionals, cybersecurity professionals, consumer protection advocates, and government officials. Use by ICANN-accredited registrars is voluntary. More information on how to opt-in to the service is available at the *Naming Services Portal for Registrars*: **https://www.icann.org/resources/pages/nsp-registrars-2018-03-26-en**

Benefits of the Service

One of the key benefits is the simplification of the request process, making it easier to identify the right registrars and provide the necessary information for efficient and timely submission and consideration of disclosure requests. Instead of filling out multiple forms with varying sets of required information, each managed by different registrars, requestors need only to complete a single, standardized form through the service.

Requestors also no longer need to look up the appropriate registrar to contact—the service will do that for them. The service also provides a centralized platform where requestors can conveniently access pending and past requests. They can create new requests, develop request templates for future use, and cancel requests when needed. Registrars can benefit from using the service as it provides a mechanism to manage and track all nonpublic data requests in a single location. Registrars can receive automated alerts anytime they receive a request. The use of a standardized submission form also makes it easier for the correct information and supporting documents to be provided to evaluate a request. For more information on the RDRS, including a flyer for requestors, visit: **https://www.icann.org/rdrs-en**

References

[1] Leslie Daigle, "WHOIS Protocol Specification," RFC 3912, September 2004.

[2] ICANN Generic Name Supporting Organization (GNSO), "Final Report of the Temporary Specification for gTLD Registration Data Phase 2 Expedited Policy Development Process," **https://gnso.icann.org/sites/default/files/file/field-file-attach/epdp-phase-2-temp-spec-gtld-registration-data-2-31jul20-en.pdf**

ADIEL AKPLOGAN is Vice President, Technical Engagement at ICANN. With more than 25 years of experience in the ICT industry (20 specifically in the Internet Technology Industry), Adiel previously served as CEO for AFRINIC (*The African Network Information Centre*), IT Director for Symbol Technology in France (2001–2003), and Director of New Technology at CAFÉ Informatique in Togo (1994–2000). He earned a graduate degree in Electrical Engineering and holds a Master's degree in E-Business and New Technology Management from Paris Graduate School of Management. Recognized as one of the Internet technology pioneers in Africa, he contributed to technical capacity building and deployment of some of the first private Internet Service Providers in Africa from 1996 to 1999. He can be reached at: **adiel.akplogan@icann.org**

Fragments

APNIC Releases Strategic Plan

The *Executive Council* (EC) of the *Asia Pacific Network Information Centre* (APNIC) is pleased to announce the availability of APNIC's new four-year strategy. The *APNIC Strategic Plan* (2024-2027)[1] was created by the APNIC EC and Secretariat. It is informed by feedback from Members and the community. The plan sets out the future that APNIC wishes to see, the objectives and priorities that need to be achieved to help reach that future state, and the guiding principles underpinning APNIC's efforts.

The existing strategic pillars of activity (Membership, Registry, Development, Information, and Capability) have been re-cast into four new ones: Two Value Streams, Registry and Development; and two Enablers, Engagement and Capability.

The EC and Secretariat believe the new strategic pillars are the best way to group APNIC's priorities and activities over the coming four years, and the Secretariat is transitioning to a new operational staffing structure to mirror the plan's four pillars.

The strategy becomes the guide for APNIC's annual *Activity Plans*[2], and the activities will align with the overall strategy. The first annual Activity Plan based on the strategy will be released in March 2024 at the APNIC 57 *Annual General Meeting* (AGM) in Bangkok, held in conjunction with APRICOT 2024: **https://2024.apricot.net/**

[1] **https://www.apnic.net/wp-content/uploads/2023/12/APNIC_Strategic_Plan_2024-27.pdf**

[2] **https://www.apnic.net/about-apnic/corporate-documents/plans-and-strategies/**

Randy Bush Receives Rob Blokzijl Award

The *2023 Rob Blokzijl Award* was presented to Randy Bush at the RIPE 87 meeting in Rome in November for his many years of contributions to the Internet in the RIPE NCC service region and beyond, playing a vital role in establishing Internet networks in many developing countries in Africa, Latin America and the Caribbean. The award committee also recognised Randy's non-technical contributions as a dedicated mentor, for speaking the truth, and for passing on knowledge and values.

The award, bestowed by the Rob Blokzijl Foundation, honours the memory of Rob Blokzijl, the first Chair of RIPE. It recognises individuals who have made substantial technical and operational contributions to the development of the Internet in the RIPE NCC service region and supported or enabled others.

You can watch the presentation here: **https://ripe87.ripe.net/archives/video/1145/**

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

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

Supporters and Sponsors

Supporters

Internet Society

CISCO

Diamond Sponsors

Your logo here!

Ruby Sponsors

APNIC FOUNDATION

ICANN

Sapphire Sponsors

Your logo here!

Emerald Sponsors

Akamai

amsix

APRICOT Asia Pacific Regional Internet Conference on Operational Technologies

.auDA

DE-CIX

EQUINIX

Google

identity digital

JPRS

JUNIPER NETWORKS

lacnic

LinkedIn

linx

netskope

NSRC Network Startup Resource Center

NTT DATA

RIPE NCC RIPE NETWORK COORDINATION CENTRE

TEAM CYMRU

VERISIGN

WIDE PROJECT

Corporate Subscriptions

AFRINIC The Internet Numbers Registry for Africa

COMCAST

DNS-OARC

Edgio

FLEXOPTIX

IPXO

ISC Internet Systems Consortium

IWL

qa cafe

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

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

CHANGE SERVICE REQUESTED

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

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