The Internet Protocol Journal

March 2024 Volume 27, Number 1

A Quarterly Technical Publication for Internet and Intranet Professionals

In This Issue

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

ISSN 1944-1134

From The Editor

I have just returned from the annual *Asia Pacific Regional Internet Conference on Operational Technologies* (APRICOT), held this year in Bangkok, Thailand. Amongst the many interesting presentations given, there was one entitled "BGP in 2023," by Geoff Huston. In his talk, he asked if we have reached "Peak IPv4," noting that the overall IPv4 routing growth trends slowed down or even reversed through 2023. His presentation, as well as a YouTube video, are available on the APRICOT 2024 website.

In our two previous issues, we published a two-part set of articles under the heading "Introduction to 5G" by William Stallings. Part One introduced the standards, specifications, and usage scenarios for 5G. Part Two gave an overview of the structure and function of 5G networks. A third article, on *Network Slicing*, which is closely related to 5G, is included in this edition.

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. Our second article, by Mikael Holmberg, describes the history and future of Ethernet.

Pindar Wong has served on our Editorial Advisory Board since the inception of this journal. I have always appreciated his invaluable insight and advice, particularly on emerging technologies such as Blockchain. Pindar has indicated that he is moving on to pursue other interests and wishes to step down from his advisory role. I thank Pindar for all his contributions and wish him the best in his future endeavors.

I am also extremely honored to welcome Merike Kaeo as a new member of the Editorial Advisory Board. Merike has extensive experience in all aspects of network and information security, and I look forward to working with her on developing article topics for IPJ.

Publication of this journal is made possible by the generous support of our donors, supporters, and sponsors. We also depend on your feedback and suggestions. If you would like to comment on, donate to, or sponsor IPJ, please contact us at **ipj@protocoljournal.org**

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

Network Slicing

by William Stallings, Independent Consultant

One of the most important features of 5G is *Network Slicing*[1,10,11]. Network slicing uses virtualization technologies, especially *Software Defined Networks* (SDN) and *Network Functions Virtualization* (NFV)[0], which enable 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, *Quality of Service* (QoS), and specific users.

The article "Network Slicing for 5G: Challenges and Opportunities,"[2] lists the following advantages of slice-based networking compared with traditional networks:

- 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 the "softwarization" techniques of 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 is achieved by (i) using a different physical resource; (ii) separating by virtualization, which may allow sharing of physical resources; or (iii) 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.

Overview

Network slicing permits a physical network to be separated into multiple virtual networks (logical segments) that can support different *Radio Access Networks* (RANs) or several types of services for certain customer segments, greatly reducing network construction costs by using communication channels more efficiently. In essence, network slicing allows the creation of multiple virtual networks atop a shared physical infrastructure.

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

The *International Telecommunication Union Telecommunication Standardization Sector* (ITU-T) is involved in the standardization of network slicing for 5G networks. ITU-T Recommendation Y.3112[3] defines a network slice as a logical network that provides specific network capabilities and network characteristics. This recommendation lays out an overall framework for network slicing, defines high-level requirements, and describes core network functions relevant to network slicing.

Figure 1 illustrates the network slicing concept. The requirements of a particular application or user determine the physical and logical network resources needed to provide the desired QoS. The network slicing function dedicates the appropriate resources to support that QoS. Figure 1 illustrates the three major usage scenarios for 5G defined by *ITU Radiocommunication Sector* (ITU-R)[4]. The scenarios include:

- *enhanced Mobile Broadband* (eMBB): Characterized by high data rates for mobile devices.
- *massive Machine-Type Communication* (mMTC): Characterized by the ability to support huge numbers of devices, such as in a large *Internet of Things* (IoT) deployment.
- *Ultra-Reliable and Low-Latency Communication* (URLLC): Characterized by the ability to support human-to-machine and machine-to-machine communications that require high reliability and/or low end-to-end delay.

Figure 1: Network Slicing Concept

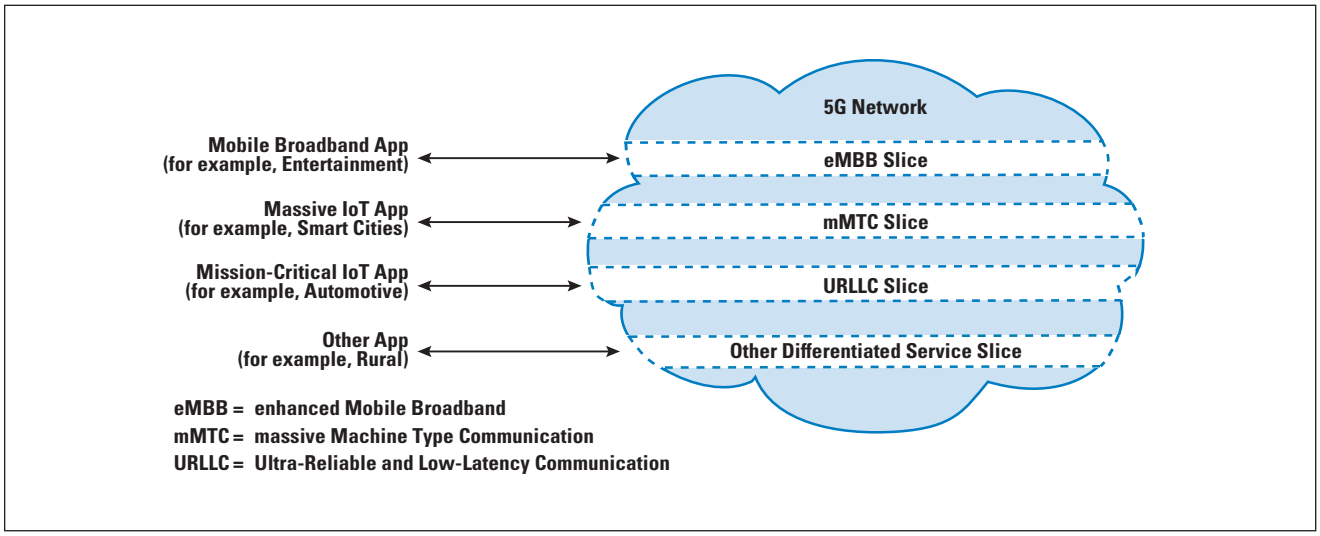

Network Slicing Concepts

Network slicing permits you to separate a physical network into multiple virtual networks (logical segments) that can support different radio access networks or several types of services for certain customer segments, greatly reducing network construction costs by using communication channels more efficiently. In essence, network slicing allows you to create multiple virtual networks atop a shared physical infrastructure. This virtualized network scenario devotes capacity to certain purposes dynamically, according to need. As needs change, so can the devoted resources. Using common resources such as storage and processors, network slicing permits you to create slices devoted to logical, self-contained, and partitioned network functions. It supports the creation of virtual networks to provide a given QoS, such as guaranteed delay, throughput, reliability, and/or priority.

A network slice creates a partition of the core network consisting of virtualized network functions and resources running on some of the core network hardware resources. Figure 2, based on concepts in a *Next Generation Mobile Networks* (NGMN) document[5], illustrates network slicing concepts. Figure 2a shows a simple core network configuration comprising three types of devices:

- *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.
- *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.
- *Access Nodes:* These nodes provide an interface to RANs, which in turn provide access to mobile *User Equipment* (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 2 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 2: 5G Network Slices Implemented on the Same Infrastructure

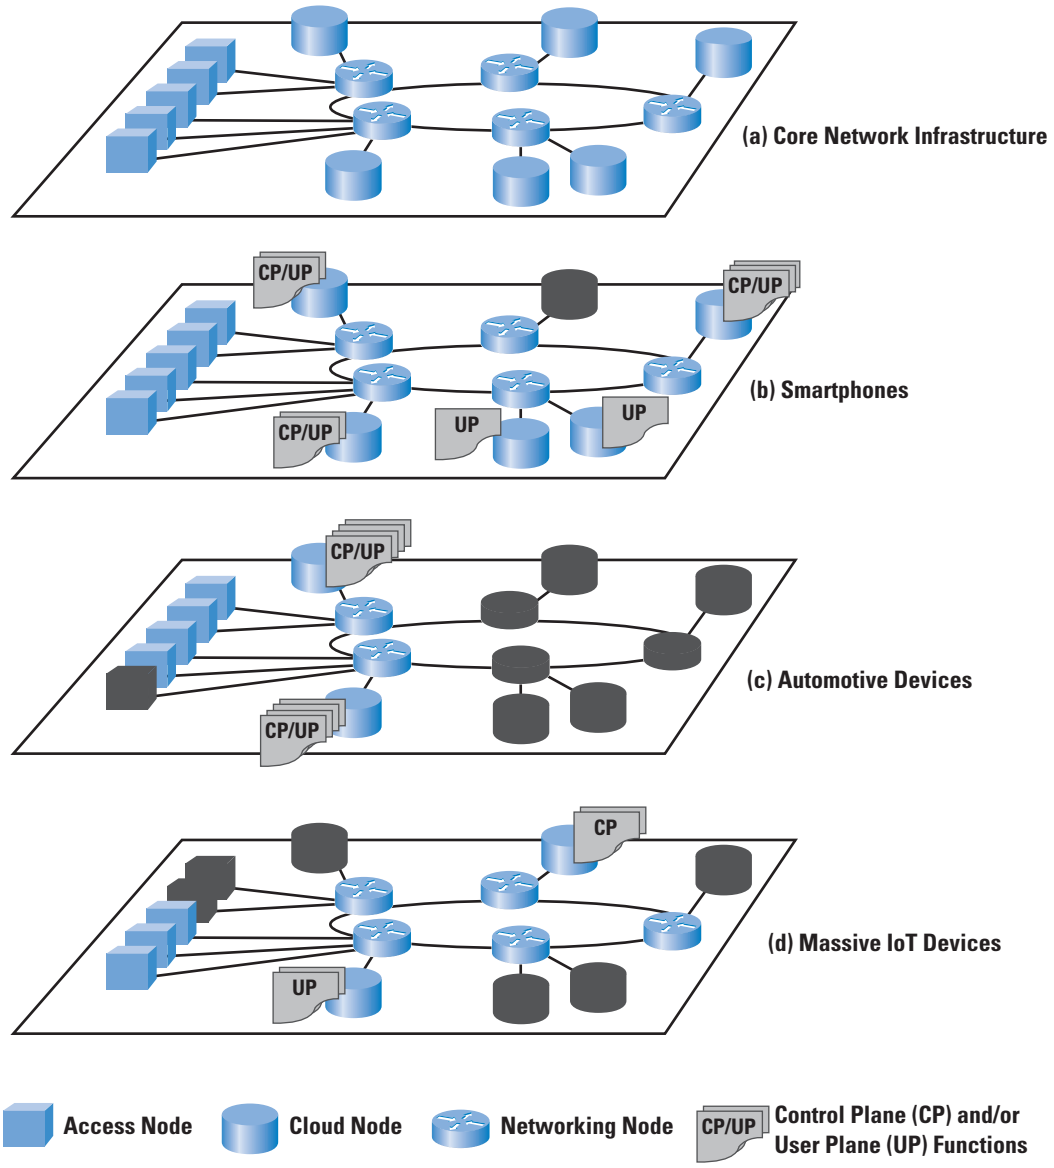

The first network slice depicted in Figure 2 is for a typical smartphone use case. Such a slice might have fully-fledged functions distributed across the network. The second network slice in Figure 2 indicates the type of support that may be allocated for automobiles in motion. This use case emphasizes the need for security, reliability, and low latency. A configuration to achieve these necessities would limit core network resources to nearby cloud-edge nodes, in addition to recruiting sufficient access nodes to support the use case.

The final use case illustrated in Figure 2 is for a massive 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.

Requirements for Network Slicing

The *3rd Generation Partnership Project* (3GPP) is the organization responsible for developing specifications that are subsequently issued as ITU-T Recommendations. The 3GPP Technical Specification TS 22.261[6] lists requirements for network slicing in two categories: *general requirements* and *management requirements*.

The general requirements for network slicing are the following:

- It must provide connectivity to home and roaming users in the same network slice.
- In a shared 5G network configuration, each operator must be able to apply all the requirements to their allocated network resources.
- It must support the *IP Multimedia Subsystem* (IMS) as part of a network slice.
- IMS support must be independent of network slices.

The IMS is a standards-based architectural framework for delivering multimedia communications services such as voice, video, and text messaging over IP networks[7,12]. 3GPP originally developed the IMS specifications in the early 2000s to standardize access to multimedia services using cellular networks. The specifications define a complete framework and architecture that enables the convergence of video, voice, data, and mobile network technologies.

The management requirements of network slicing follow; it must:

- Allow the operator to create, modify, and delete a network slice.
- Allow the operator to define and update the set of services and capabilities supported in a network slice.
- Allow the operator to configure the information that associates a UE to a network slice.
- Allow the operator to configure the information which associates a service to a network slice.
- Allow the operator to assign a UE to a network slice, to move a UE from one network slice to another, and to remove a UE from a network slice based on subscription, UE capabilities, the access technology the UE uses, and the operator's policies and services the network slice provides.

- Support a mechanism for the *Visited Public Land Mobile Network* (VPLMN), as authorized by the *Home Public Land Mobile Network* (HPLMN), to assign a UE to a network slice with the needed services or to a default network slice.
- Enable a UE to be simultaneously assigned to and access services from more than one network slice of one operator.
- Ensure traffic and services in one network slice will have no impact on traffic and services in other network slices in the same network.
- Ensure that the creation, modification, and deletion of a network slice will have no or minimal impact on traffic and services in other network slices in the same network.
- Support scaling of a network slice, that is, adaptation of its capacity.
- Enable the network operator to define a minimum available capacity for a network slice. Ensure that scaling of other network slices on the same network will have no impact on the availability of the minimum capacity for that network slice.
- Enable the network operator to define a maximum capacity for a network slice.
- Enable the network operator to define a priority order between different network slices in case multiple network slices compete for resources on the same network.
- Support means by which the operator can differentiate policy control, functionality, and performance provided in different network slices.

Identification and Selection of a Network Slice

The *Single Network Slice Selection Assistance Information* (S-NSSAI) defines a single network slice. An S-NSSAI consists of two elements:

- *Slice/Service Type* (SST): An identifier that refers to the expected slice behavior in terms of features and services. Standardized SST values provide a way for establishing global interoperability for slicing so that 5G networks can support the roaming use case more efficiently for the most commonly used SSTs. Table 1 lists the standardized SSTs.
- *Slice Differentiator* (SD): Optional information that complements the SST to differentiate among multiple network slices of the same SST.

Table 1: Standardized Slice/Service Type Values.

Slice/Service Type	SST Value	Characteristics
eMBB	1	Slice suitable to handle 5G-enhanced Mobile Broadband.
URLLC	2	Slice suitable to handle ultra-reliable low-latency communications.
Massive IoT (mIoT)	3	Slice suitable to handle massive IoT.
Vehicle-to-Everything (V2X)	4	Slice suitable to handle V2X services.

A UE may be served by up to eight network slices at a time, each identified by an S-NSSAI. The set of S-NSSAIs associated with a UE form a *Network Slice Selection Assistance Information* (NSSAI) data object.

Functional Aspects of Network Slicing

Figure 3 indicates the manner in which core NFs are used to implement network slices. A network function is a processing function in a network that has defined functional behavior and interfaces. You can implement a network function as a network element on dedicated hardware, as a software instance running on dedicated hardware, or as a virtualized function instantiated on an appropriate platform, for example, on a cloud infrastructure. Some NF instances support multiple network slices serving a UE, while others are specific to a given slice.

Figure 3: Network Functions that Support Network Slicing

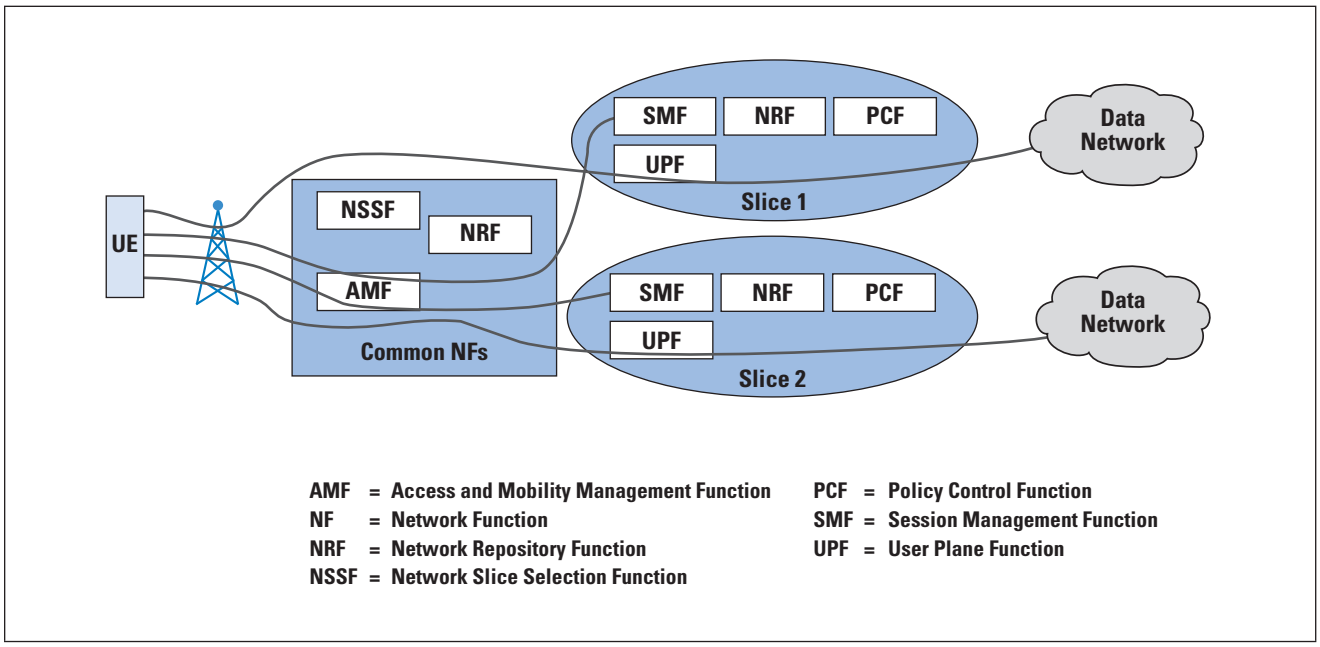

The common NFs follow:

- *Access and Mobility Management Function* (AMF): Network slice instance selection is usually triggered as part of the registration procedure by the first AMF that receives the registration request from the UE. When a UE accesses the network, AMF provides functionalities to register and de-register the UE with the network, and it establishes the user context in the network. In the registration procedure, AMF performs, but is not limited to, network slice instance selection, UE authentication, authorization of network access and network services, and network access policy control. In addition, when AMF receives a session establishment request message from UE, it performs discovery and selection of the SMF that is the most appropriate to manage the session.

- *Network Slice Selection Function* (NSSF): The AMF retrieves the slices that the user subscription allows and interacts with the NSSF to select the appropriate network slice instance (for example, based on allowed S-NSSAIs, 5G network ID, and other parameters). The NSSF responds with a message including the list of appropriate network slice instances for the UE. As a result, the registration process may switch to another AMF if needed.
- *Network Repository Function* (NRF): During the AMF-NSSF interaction, the NSSF may return the identity of one or more NRFs to be used to select NFs and services within the selected network slice instance(s).

The slice-specific NFs follows:

- *Session Management Function* (SMF): The UE sends a message to the AMF requesting that a *Protocol Data Unit* (PDU) session be associated to one S-NSSAI and one *Data Network* (DN). The AMF selects the appropriate SMF, which manages the PDU session. The SMF sets up the PDU session for the UE and controls the user-plane operation. The SMF selects the UPF and invokes enforcement of QoS and charging policies.
- *User Plane Function* (UPF): Once a PDU session is established, QoS flows for this PDU session over this network slice pass through the UPF.
- *Policy Control Function* (PCF): The SMF gets policy information related to session establishment from the PCF.
- *Network Repository Function* (NRF): The SMF uses the NRF to discover the required NFs for the individual network slice.

Generic Slice Template

3GPP TS 28.531[8] includes a description of the concept of the *Generic Slice Template* (GST). The *GSM Association* (GSMA) has specified the GST, which provides a standardized list of attributes that you can use to characterize different types of network slices[9]. A *Network Slice Type* (NEST) is a GST filled with (ranges of) values. There may be two kinds of NESTs:

- *Standardized NEST* (S-NEST): Attributes are assigned (ranges of) values by *Standards-Developing Organizations* (SDOs), working groups, forums, and so forth, such as 3GPP, GSMA, *5G Automotive Association* (5GAA), and the *5G Alliance for Connected Industry and Automation* (5G-ACIA).
- *Private NEST* (P-NEST): Attributes are assigned (ranges of) values by the Network Slice Providers; these values are different from those assigned in S-NESTs.

Network Slice Providers can build their network slice product offering based on S-NESTs and/or their P-NESTs. GSMA has developed the GST to be a list of attributes sufficient for describing a wide range of NESTs that you can fully construct by allocating values (or ranges of values) to each relevant attribute in the GST. A network operator can use a NEST to identify the network resources and functions needed to instantiate network slices. The process to fill in the GST and to create a NEST comprises three steps:

1. Study use cases and derive service requirements based on discussions with the slice customers, such as vertical industries or specific enterprises.
2. Convert the service requirements identified in (1) into technical requirements.
3. Document the technical requirements produced in (2) using the NEST by filling in the values of each of the attributes of the GST.

The current version of the GST lists 35 attributes, shown in Figure 4.

Figure 4: Generic Network Slice Template Attributes

Availability	Network Functions Owned by Network Slice Customer	Supported Device Velocity
Area of Service	Maximum Number of PDU Sessions	Synchronicity
Delay Tolerance	Maximum Number of UEs	UE Density
Deterministic Communication	Performance Monitoring	Uplink Throughput per Network Slice
Downlink Throughput per Network Slice	Performance Prediction	Uplink Maximum Throughput per UE
Downlink Maximum Throughput per UE	Positioning Support	User Management Openness
Energy Efficiency	Radio Spectrum	User Data Access
Group Communication Support	Root Cause Investigation	V2X Communication Mode
Isolation Level	Session and Service Continuity Support	Latency from (last) User Plane Function (UPF) to Application Server
Maximum Supported Packet Size	Simultaneous Use of the Network Slice	Network Slice Specific Authentication and Authorization (NSSAA) Required
Mission-Critical Support	Slice Quality of Service Parameters	
Multimedia Telephony (MMTel) Support	Support for Non-IP Traffic	
NB-IoT Support		

Summary

Virtualization encompasses a variety of technologies for managing computing resources by providing a software translation layer, known as an abstraction layer, between the software and the physical hardware. Virtualization turns physical resources into logical, or virtual, resources. Virtualization enables users, applications, and management software operating above the abstraction layer to manage and use resources without needing to be aware of the physical details of the underlying resources. NFV is a key technology for implementing 5G wireless networks.

References

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

[1] William Stallings, *5G Wireless: A Comprehensive Introduction*, ISBN-13: 9780136767299, Pearson, 2021.

[2] 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.

[3] ITU-T, "Framework for the support of network slicing in the IMT-2020 network," ITU-T Y.3112, December 2018

[4] ITU-R, "IMT Vision—Framework and overall objectives of the future development of IMT for 2020 and beyond," ITU-R Recommendation M.2083, September 2015.

[5] Next Generation Mobile Network Alliance, "5G End-to-End Architecture Framework," Version 4.31, November 2020.

[6] 3GPP TS 22.261, "Technical Specification Group Services and System Aspects; Service requirements for the 5G system; Stage 1 (Release 18)," January 2021.

[7] Martin Koukal and Robert Bestak, "Architecture of IP Multimedia Subsystem," Proceedings ELMAR Symposium, June 2006.

[8] 3GPP TS 22.531, "Technical Specification Group Services and System Aspects; Management and Orchestration; Provisioning; (Release 16)," April 2020.

[9] GSM Association, "Generic Network Slice Template Version 3.0," May 22, 2020.

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

[11] William Stallings, "Introduction to 5G Part Two: Core Network, Radio Access Network, and Air Interface," *The Internet Protocol Journal*, Volume 26, No. 3, December 2023.

[12] Mark Grayson, "Lessons Learned from 20 Years of Cellular and Wi-Fi Integration," *The Internet Protocol Journal*, Volume 26, No. 3, December 2023.

WILLIAM STALLINGS is an independent consultant and author of numerous books on computer networking, security, and computer architecture. His latest book is *Wireless 5G: A Comprehensive Introduction* (Pearson, 2021). He maintains a computer science resource site for computer science students and professionals at **ComputerScienceStudent.com** and is on the editorial board of *Cryptologia*. He has a Ph.D. in computer science from M.I.T. and can be reached at **wllmst@icloud.com**

The History and Future of Ethernet

by Mikael Holmberg, Extreme Networks

Initially developed by *Xerox Palo Alto Research Center* (PARC) in the 1970s and ratified by the *Institute of Electrical and Electronics Engineers* (IEEE) as a standard in 1983, the evolution of Ethernet has taken this technology through many specifications and standardizations during its 50-year history.

Ethernet technology has become the backbone of modern communication and connectivity, linking billions of devices to each other and the Internet. Today, Ethernet connects *Local Area Networks* (LANs), *Wide Area Networks* (WANs), Internet, Cloud, *Internet of Things* (IoT) devices, Wi-Fi, and many other systems into one seamless global communications network.

The name *Ethernet* is based on the word "ether" as a way of describing an essential feature of the system: the physical medium (that is, a cable) carries bits to all stations, much the same way that the old "luminiferous ether" was once thought to propagate electromagnetic waves through space.

In its early days, Ethernet competed with other technologies like *Token Ring*. It was eventually chosen as the ubiquitous technology used in computer networks because of the simplicity by which the communication protocol can be deployed and its ability to incorporate modern advancements without losing backward compatibility. Ethernet continues to reign as the de facto standard for computer networking and many newly evolved applications and use cases. Just to choose one of interest, the topic that everybody talks about today is *Artificial Intelligence* (AI). As AI workloads increase, network industry giants are teaming to ensure Ethernet networks can keep pace and satisfy the AI high-performance networking requirements, among many other new use cases and applications. I will cover AI as well as a few other interesting use cases around the evolved Ethernet in this article.

In 1975, Xerox filed a patent application listing Bob Metcalfe, David Boggs, Chuck Tucker, and Butler Lampson as inventors. Then, in 1976, after the system was deployed at PARC, Metcalfe and Boggs published a seminal paper.[0] Four gentlemen, Yogen Dalal, Ron Crane, Bob Garner, and Roy Ogus, facilitated the upgrade from the original 2.94-Mbps to the 10-Mbps protocol, which was released to the market in 1980 and ratified by the IEEE as a standard in 1983[2]. Ethernet has become the dominant LAN technology, and five decades after its initial specification its evolution continues.[31]

Taking a step back in time, let's look at the progress of Ethernet technology over the past five decades and explore where experts think it could be heading in the years to come.

The Early Days of Ethernet

The evolution of Ethernet officially began in 1973 when engineer Robert Metcalfe introduced the concept in a memo he wrote while working at Xerox PARC. Metcalfe initially described Ethernet as interconnecting computing workstations, and it enabled them to communicate with each other as well as with devices like laser printers. These interconnected endpoints became the environment we now recognize as the world's first LAN.

Metcalfe was inspired by ALOHAnet[1], an earlier networking project that began at the University of Hawaii in 1968 and aimed to connect remote workstations across the Hawaiian Islands to a central computer at the main Oahu campus.

ALOHAnet was realized by using a quite rudimentary *Additive Links On-line Hawaii Area* (ALOHA) protocol, where an end station would transmit a frame over a common data channel and then wait for confirmation that it had reached its destination successfully. If the end station didn't receive confirmation within a given period, it assumed a *collision* had occurred with another frame sent by a different end station simultaneously. In that case, that station would continue to resend the data until it achieved successful transmission. But as amounts of end stations and transmissions increased, more collisions would occur, and the network would become less efficient. An ALOHA variation named *Slotted ALOHA* aimed to minimize network contention problems by precisely coordinating individual transmissions for the end stations and assigning them designated timeslots via a beacon signal schema.

Metcalfe's Ethernet experiment, at that time referred to as the *Alto ALOHA Network*, included many revolutionary features that enabled significantly more efficient use of a computer network. This set of rules, which became known as the *Carrier Sense Multiple Access/Collision Detect* (CSMA/CD) protocol, allowed end stations to monitor the availability of a shared communication path and detect possible collisions when two end stations sent data at the same time. When frames collided, the system would discard them, leading each end station to wait for a randomly assigned length of time before trying to resend. The end station would continue this schema to pause and try to resend as many times as necessary. This process is known as *exponential back-off*.[33]

So, the original Ethernet technology was based on a shared medium that was collision-prone, where all computers trying to communicate shared the same cable and, as such, competed with each other. The modern Ethernet implementation has a collision-free switched connection, where each computer communicates with only its own switch port, without competing for the cable with others.

By 1973, Metcalfe thought that the technology had outgrown its original name and renamed it *Ethernet*. Four years later, Metcalfe and Boggs, together with Charles Thacker and Butler Lampson, who also worked at Xerox, successfully patented Ethernet technology.[3]

How Does Ethernet Work?

Ethernet works by breaking up data being sent to or from devices, like a personal computer, into short pieces of different-sized bits of information called *frames*. Those frames contain information such as the source and destination address that helps the frame route its way through a network.

In the past, computers on a LAN typically shared a single connection. Ethernet was built around the principle of CSMA/CD, as was briefly explained earlier in this article, where the protocol ensures that the cable is not in use before sending any frames out. Now that feature is far less important than it was in the early days of networking, as devices generally have their own private connection to a network through a *switch*. Ethernet now operates using *Full Duplex* (FDX), where the sending and receiving channels are separate, so it is impossible for collisions to occur over the same connection. As there is no error correction in Ethernet, the communication relies on upper-layer advanced protocols to ensure that everything is transmitted perfectly. Ethernet provides the basis for most digital communications, and it integrates quite easily with most higher-level protocols.

Ethernet IEEE Standardization

Xerox worked with two other vendors, Digital Equipment Corporation and Intel, to publish the first 10-Mbps Ethernet specification in 1980. Meanwhile, the *Local and Metropolitan Area Networks* (LAN/MAN) Standards Committee at the IEEE set out to develop a similar open standard. The IEEE LAN/MAN committee, which applies the number 802 to all its standards, formed an Ethernet subcommittee and named it the *IEEE 802.3 Working Group*. Through the first half of the 1980s, the Ethernet 10BASE-5 implementation used a coaxial cable 0.375 inches (9.5 mm) in diameter, also referred to as *Thick Ethernet* or *Thicknet*, and it was standardized in 1982 as 10BASE-5.

Figure 1: 10BASE-5 Ethernet

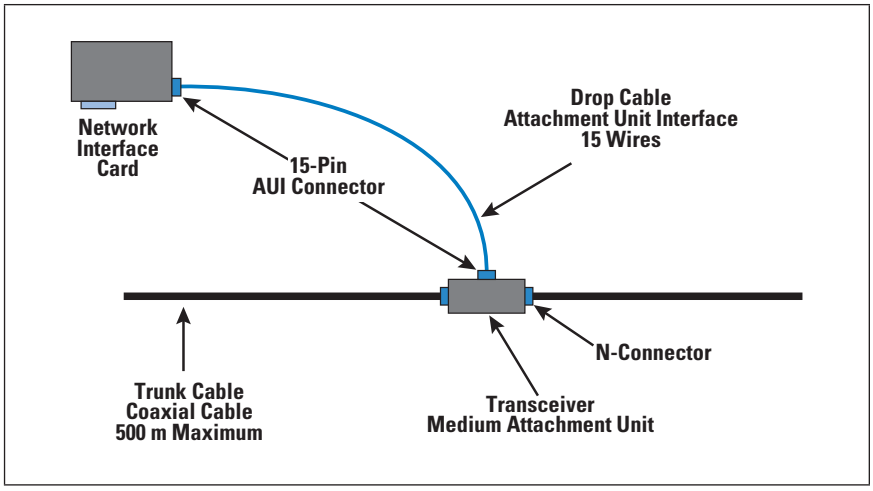

Some of us who have worked in this industry for many years might recall that on 10BASE-5, you drilled the transceiver pin (the so-called "vampire tap") into the core of the thick coaxial cable, and if you were not careful, you might end up shortcutting the wire.

In the late 1980s, 10BASE-5 was replaced by 10BASE-2, called *Thin Ethernet* or *Thinnet*, and it used a BNC connector to connect the Ethernet *Network Interface Cards* (NICs) to a BNC-T splitter to ensure that the Ethernet segment stayed intact. It used the RG-58 coaxial cable, which is 0.2 inches (5 mm) in diameter, as media. The emphasis was on making the installation of the cable easier and less costly. This thin Ethernet was followed by twisted pair (10BASE-T) and fiber-optic (10BASE-FL).

Figure 2: 10BASE-2 Ethernet

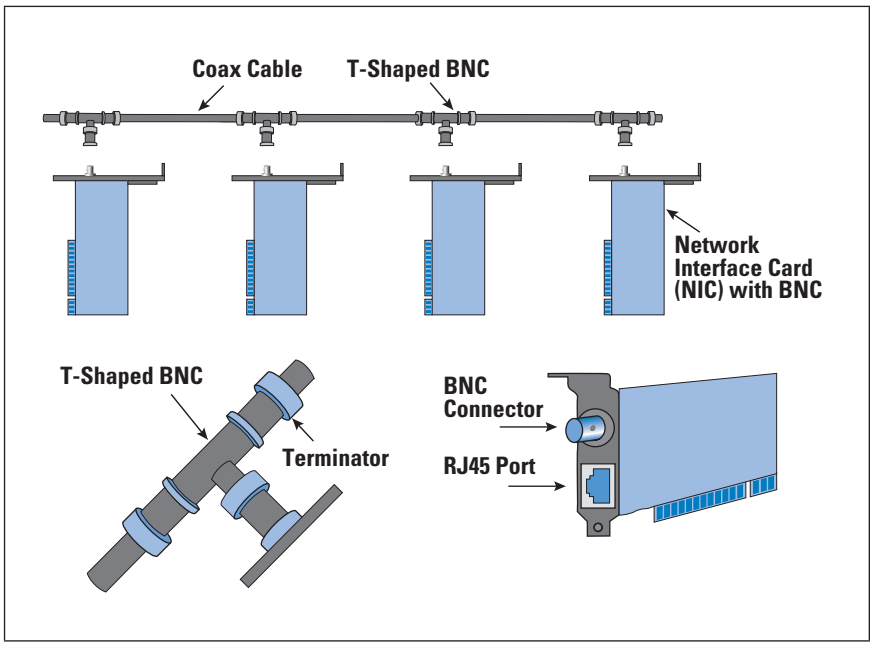

In 1995, with the Fast Ethernet standard, the speed was upgraded to 100 Mbps, and no such speed improvement was ever made for Thinnet. By 2001, prices for Fast Ethernet cards had fallen to under $50, and by 2003 Wi-Fi (802.11) networking equipment was widely available and affordable. Because of the immense demand for high-speed networking, the low cost of *Category 5* (Cat5) cable, and the popularity of 802.11 wireless networks, both 10BASE-2 and 10BASE-5 have become obsolete, though devices using those standards might still exist in some locations.

Also, in 1995, 100-Mbps Ethernet introduced *auto-negotiation*, which allowed for two network devices to signal each other and establish the best-shared mode of operation, including speed and duplex mode.

Three years later, a new milestone was reached when the 802.3 working group introduced *Gigabit Ethernet* [GE] (100BASE-T)[4, 5], which was first realized over fiber-optic cable and, subsequently, over twisted-pair copper cable.

The evolution of Ethernet continued with 10-Gbps speeds in 2002, initially over fiber, then twinax cable, and finally, over unshielded twisted-pair cable. Then, in 2010, IEEE approved 40 GE and 100 GE, which was realized by aggregating multiple 10-Gbps lanes.

In 2016, driven by the rising demand from hyperscalers (web companies), the IEEE ratified 25 GE, which was 2.5 times faster than 10 GE but more cost-efficient than 40 GE. This standard improved throughput by increasing the capacity of a single lane, rather than aggregating multiple lower-capacity lanes, and meant that 25 GE required less cable and power and had higher port density than 40 GE. In some cases, an upgrade to 25 GE lets data-center operators extend the life of top-of-rack switches and avoid full "rip-and-replace" upgrades of cabling infrastructure. Hence, hyperscalers upgraded to 25 GE speeds in their data centers.

The following year, in late 2017, the networking industry saw the ratification of 200 and 400 GE. These standards were both based on 50-Gbps single lanes, as the cloud providers and hyperscale data centers, *Internet Service Providers* (ISPs), and specialized organizations like *Network Operations Centers* (NOCs) needed and wanted more bandwidth. Some of the challenges with 400-Gbps speeds include new cabling requirements because the current Category 5 and 5e cables don't support such speeds.

In 2019, *Communication Service Providers* (CSPs) began deploying (or, more likely, testing) 5G [26, 27] networks, the fifth-generation technology standard for broadband cellular networks. It is defined by the *3rd Generation Partnership Project* (3GPP), and it is the planned successor to 4G networks. Like its predecessors, 5G networks are cellular networks. All 5G wireless devices in a cell connect to the Internet and telephone network via radio waves through a local antenna in the cell. These new networks boost higher download speeds, eventually up to 10 Gbps. In addition to being faster than existing networks, 5G offers higher bandwidth, enabling it to connect a greater number of devices and improve the quality of Internet services in crowded areas. Naturally, Ethernet acts as the packet-based solution within 5G, accommodating all the essential containerized microservices required for 5G functions running on computers in all sizes of data centers with Ethernet fabric technologies.

Ethernet-based 5G cloud data-center fabrics come in various sizes, from small edge data-center fabrics implemented as Layer 2 network infrastructures to truly scalable three- and five-stage large data-center fabrics. These larger fabrics deployed as Layer 3 infrastructure with dozens or even hundreds of Ethernet switches connected in a *spine and leaf* architecture, also known as CLOS. The CLOS architecture has its origins in Charles Clos' crossbar switches for telephone-call switching, and it is composed of leaf and spine layers where switches are used.

The most prevalent design for these cloud data-center fabrics consists of Ethernet switches that use *Virtual Extensible LAN* (VXLAN) with *Multiprotocol Extensions for BGP* (MP-BGP) and an *Ethernet Virtual Private Network* (EVPN) control plane.

All Ethernet switches are deployed in pairs to provide dual-homed redundant connectivity to computers and other switches. The leaf switch pairs interconnect to form a cluster, providing redundancy for the attached computers. *Border Leaf* (BL) switches are also deployed in pairs, ensuring dual-homed redundant connectivity to external *Provider Edge* (PE) routers and the Internet. This connectivity presents yet another interesting use case in which Ethernet serves as the foundation for cloud-native 5G mobile network applications and workloads.

Technically, the specification for 800-Gbps Ethernet also exists but is not really used outside of test environments. The interesting thing about Ethernet is that because it is such an open protocol, there is no reason to think that even the 800-Gbps speeds are anywhere near the theoretical maximum. Research is being done to set the groundwork for a 1.6-Tbps standard. Speeds like that will probably be useful only in highly specific applications.

Ethernet Cables

You can't talk about Ethernet without also talking about various cables used for Ethernet. As I previously described, the early days of Ethernet relied on coaxial cables, basically the same as were used for cable television. Coaxial cable is robust in design, having a thick internal copper wire, but it does have trade-offs, as it is heavy and difficult to work with and not very flexible. Ethernet changed to use twisted-pair cables that are still used when deploying Ethernet networks today, as are fiber-optic cables.

Many companies that manufacture Ethernet cables have moved away from the dull gray color scheme and instead offer them in a wide variety of colors that allow for improving data-center racks with different-colored cables. It also enables color-coding, so technicians can group their different network connections visually into groups based on different colors for quick troubleshooting.

The standard plug on both ends of twisted-pair cables (RJ-45), which is very similar to the same kind of connector that wired telephone systems use (RJ-11), made it easy to just click the cables into any device that supports Ethernet connectivity. So, simply plugging in a device and attaching it to a network using one of those colored Ethernet cables is the only step required to gain connectivity. The long-time standard for Ethernet cables is *Category 5* (Cat5). The Cat5 standard has been used since 2001, and a slightly more advanced cable called *Category 5e* (Cat5e) is also used today for faster Ethernet applications. Category 5e cables are targeted at 100-Mbps Ethernet, but the design also supports higher speeds, such as Gigabit Ethernet. New *Category 6* (Cat6) cables have been introduced to support higher speeds than the Cat5 and Cat5e cables.

Several standards have been defined for *Power over Ethernet* (PoE), which allows you to connect devices with a single Ethernet cable without the need for additional power sources. Table 1 lists Ethernet cabling standards.

Table 1: Ethernet Cabling Standards

ETHERNET TYPE	BANDWIDTH	CABLE TYPE	MAXIMUM DISTANCE
10BASE-T	10Mbps	Cat 3/Cat 5 UTP	100m
100BASE-TX	100Mbps	Cat 5 UTP	100m
100BASE-TX	200Mbps	Cat 5 UTP	100m
100BASE-FX	100Mbps	Multi-mode fiber	400m
100BASE-FX	200Mbps	Multi-mode fiber	2Km
1000BASE-T	1Gbps	Cat 5e UTP	100m
1000BASE-TX	1Gbps	Cat 6 UTP	100m
1000BASE-SX	1Gbps	Multi-mode fiber	550m
1000BASE-LX	1Gbps	Single-mode fiber	2Km
10GBASE-T	10Gbps	Cat 6a/Cat 7 UTP	100m
10GBASE-LX	10Gbps	Multi-mode fiber	100m
10GBASE-LX	10Gbps	Single-mode fiber	10Km

Figure 3: Ethernet Switching, Wi-Fi, and PoE

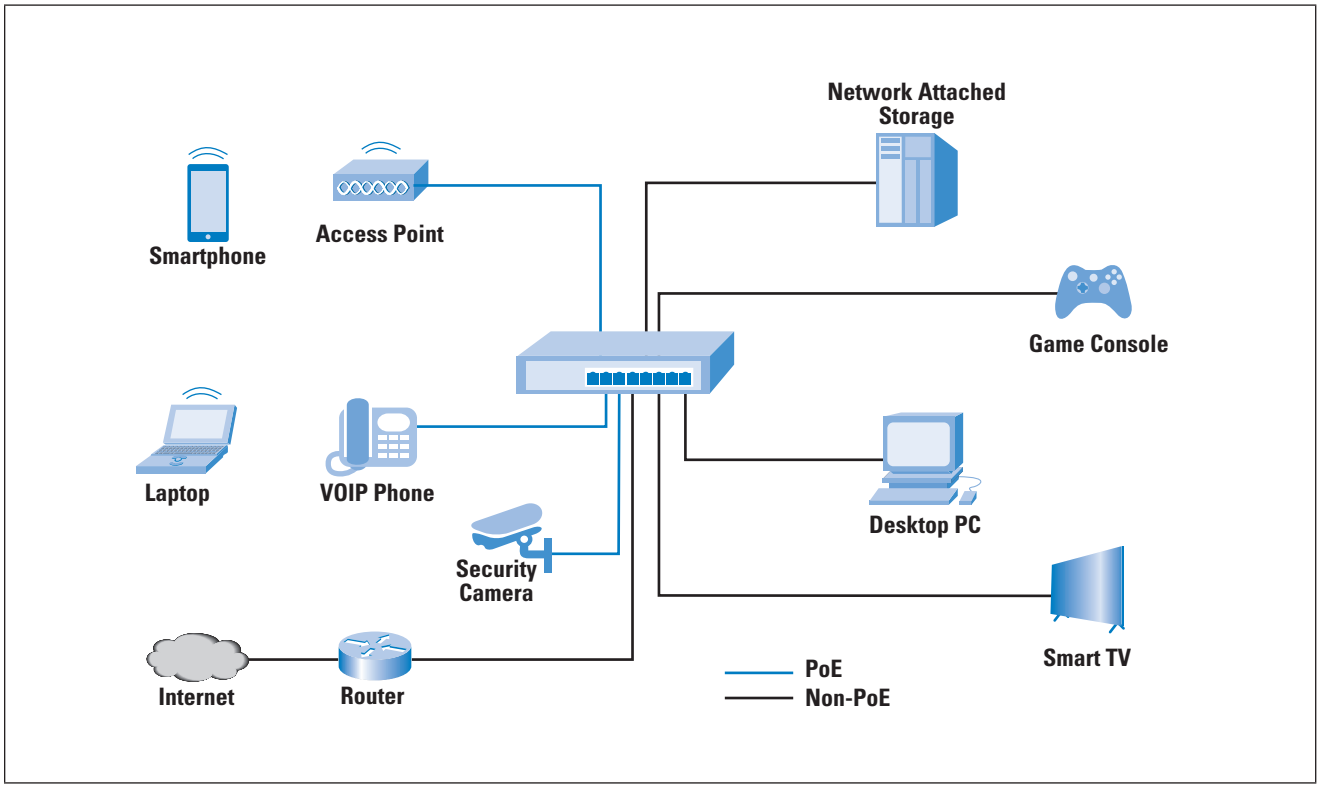

Ethernet and Time Synchronization – Are We in Sync?

As applications continue to advance, latency has become a significant concern that we need to address. The solution to this problem lies in using *Precision Time Protocol* (PTP)[28, 29], as we are addressing timing accuracy in the range of hundreds of nanoseconds. The use of Ethernet in mission-critical networks and the telecom industry showcases that Ethernet has now emerged as the de facto transport technology, using protocols like PTP to synchronize clocks throughout the network.

To appreciate the significance of PTP, it's important to understand that we're addressing timing accuracy in a range of hundreds of nanoseconds. This problem represents extremely tight timing requirements for certain applications and use cases. Maintaining precise timing is crucial for operating distributed systems at scale while ensuring that various operations remain synchronous. Additionally, precise timing is especially crucial when handling critical processes that govern infrastructure operations.

An excellent example of such systems can be seen in the telecom industry's 5G networks. Coordinating time between multiple servers or base stations in 5G could be compared to synchronized swimming in the Olympics, where all swimmers must perform their part of the routine at the same pace. If they perform at different paces, the routine will not look as it should. Ensuring that all servers and base stations operate in sync is vital for efficient network performance, much like how synchronized swimming relies on all swimmers to perform their parts of the routine at the same pace.

In addition to managing latency, new use cases for Ethernet also require it to become a deterministic networking technology, and as I briefly discussed 5G and PTP previously, *Time-Sensitive Networking* (TSN) comes into play, where PTP is one of the requirements for Ethernet to become a deterministic networking technology.

Ethernet and TSN for Deterministic Networking

Ethernet has evolved to incorporate various applications and technologies. TSN enables the synchronization of network elements and endpoints, such as switches and routers, to prioritize traffic classes and provide accountable delay and guaranteed bandwidth reservation. TSN is based on numerous international standards that are integrated with the Ethernet standard IEEE 802.3, where punctuality is ensured, allowing for transmission within a given period while simultaneously accommodating a mix of other communication protocols.

When discussing use cases and Ethernet, I need to mention *Industry 4.0* needs as the fourth industrial revolution that is transforming the manufacturing industry towards more efficient, connected, and flexible factories of the future. With Industry 4.0, factories will be able to rely on cloud-native technologies and connectivity based on Ethernet and TSN.

The goal of Industry 4.0 is to create full transparency across all processes and assets at all times, including both the *Information Technology* (IT) and *Operational Technology* (OT) domains based on architectures that require communication between production systems, logistics chains, people, and processes to unite these two domains into a single domain.

Ethernet has been used as the wired solution in both computer and automation networks. Ethernet open standard allows you to connect end devices quickly and simply as well as easily scale them to exchange data with other devices and functions. Ethernet was not originally designed to meet the requirements set by automation technology regarding guaranteed and real-time communication. Various bus systems in automation have evolved using Ethernet on a physical level while implementing proprietary real-time protocols such as PROFIBUS, PROFINET, and EtherCAT, to name a few. These systems often lead to the exclusive use of the network infrastructure as well as vendor dependencies. Today such networks handling time-critical data traffic are separated from networks directing less-critical data traffic. In the future, Industry 4.0 applications will require increasingly more consistent Ethernet networks, and TSN will provide a solution for that.

Traditional Ethernet networks involving sectors like manufacturing are based on a hierarchical automation model that separates information technology from operational technology. The IT domain includes enterprise-like communication with typical end devices such as computers, while the OT domain includes systems, machines, and software used for process control and automation. These two areas are fundamentally different in how they communicate, where the IT domain requires bandwidth while the OT domain requires high availability. Data traffic in the IT domain can be classified more as non-critical, while data traffic in the OT domain is critical and time-sensitive, and as such each domain uses a particular communication standard. Ethernet, as we know it in the enterprise or IT domain, relies on TCP/IP, while the OT domain relies on various bus systems, also known as *fieldbus systems*.

In the IT domain today, wireless technologies like Wi-Fi are used and could be used in some parts of the OT domain, but because of the nature of the technology, which is based on unlicensed spectra, it cannot guarantee bounded low latency with high reliability as the load increases. In certain use cases, Wi-Fi does not perform that well during uncontrolled interference because it uses an unlicensed spectrum. That may not be relevant for less-critical applications, because there will be a variety of applications with different traffic profile demands in both the Ethernet-enabled IT and OT domains within Industry 4.0.

Future of Ethernet

Our new world of AI workloads is expected to put unprecedented performance and capacity demands on networks that are based on Ethernet. Hence, we are possibly looking at a new enhancement to the well-known Ethernet technology to handle the scale and speed required by AI.

A group of vendors and operators have teamed to form the *Ultra Ethernet Consortium* (UEC)[30], as there are concerns that today's traditional network interconnects cannot provide the required performance, scale, and bandwidth to keep up with AI demands. The consortium aims to address these concerns by adding new capabilities to the known and proven Ethernet technology specification, adding numerous new features and capabilities including:

- Multipathing and packet spraying to ensure AI workflows have access to a destination simultaneously.
- Flexible delivery order to make sure Ethernet links are optimally load-balanced while ordering is enforced only when the AI workload requires it in bandwidth-intensive operations.
- Congestion control mechanisms to ensure AI workloads avoid hotspots and spread the load evenly across multipaths within the network. These mechanisms can be designed to work in conjunction with multipath packet spraying, thus enabling a reliable transport of AI traffic.
- End-to-end telemetry to manage congestion, where information originating from the network can advise the participants of the location and cause of the congestion. In addition, shortening the congestion signalling path and providing more information to the endpoints allow more responsive congestion control.

After this journey covering some highlights after Ethernet has enjoyed five decades of existence, one might contend that Ethernet is one of the most crucial technologies today, even though it often goes unnoticed. Ethernet, as the ubiquitous network technology, powers infrastructure across the cosmos as it is used in space as well as in the deepest ocean trenches.

I named just a few examples, including the new era of cloud-native 5G data centers that provide the infrastructure for 5G applications and workloads, the industry revolution (Industry 4.0), as well as the AI challenges. But as we all know, the number of applications being developed that have substantial requirements not only around bandwidth but also latency, etc. is increasing.

This demand requires that the underlying Ethernet transport technologies can cater to such requirements; consequently, 400 GE is a reality today, and 800 GE is expected to become commonplace in the near future. Given this trend, it wouldn't be surprising to see 1-Terabyte Ethernet in use by 2030.

Based on our unending appetite for bandwidth, Ethernet, a 50-year-old technology, will soon reinvent itself once more.

References and Further Reading

[0] Robert Metcalfe and David Boggs, "Ethernet: Distributed Packet Switching for Local Computer Networks," *Communications of the ACM*, Volume 19, Issue 7, July 1976.

[1] Norman Abramson, "The ALOHA System – Another Alternative for Computer Communications," *Proceedings of the 1970 Fall Joint Computer Conference, American Federation of Information Processing Societies* (AFIPS), Houston, Texas, November 17–19, 1970.

[2] John Shoch, Yogen K. Dalal, David D. Redell, and Ronald C. Crane, "Evolution of the Ethernet Local Computer Network," *Computer*, August 1982.

[3] Robert M. Metcalfe, David R. Boggs, Charles P. Thacker, and Butler W. Lampson, "Multipoint Data Communication System with Collision Detection," United States Patent US4063220A, December 13, 1977.

[4] William Stallings, "Gigabit Ethernet," *The Internet Protocol Journal*, Volume 2, No. 3, September 1999.

[5] William Stallings, "Gigabit Ethernet: From 1 to 100 Gbps and Beyond," *The Internet Protocol Journal*, Volume 18, No. 1, March 2015.

[6] Howard Frazier and Howard Johnson, "Gigabit Ethernet: From 100 to 1,000 Mbps," *IEEE Internet Computing*, January/February 1999.

[7] Serag Gadelrab, "10-Gigabit Ethernet Connectivity for Computer Servers," *IEEE Micro*, May-June 2007.

[8] Shamus McGillicuddy, "40 Gigabit Ethernet: The Migration Begins," *Network Evolution E-Zine*, December 2012.

[9] Gautam Chanda and Yinglin (Frank) Yang, "40 GbE: What, Why & Its Market Potential," *Ethernet Alliance White Paper*, November 2010.

[10] Mark Nowell, Vijay Vusirikala, and Robert Hays, "Overview of Requirements and Applications for 40 Gigabit and 100 Gigabit Ethernet," *Ethernet Alliance White Paper*, August 2007.

[11] John D'Ambrosia, David Law, and Mark Nowell, "40 Gigabit Ethernet and 100 Gigabit Ethernet Technology Overview," *Ethernet Alliance White Paper*, November 2008.

[12] Hidehiro Toyoda, Goichi Ono, and Shinji Nishimura, "100GbE PHY and MAC Layer Implementation," *IEEE Communications Magazine*, Volume 48, Issue 3, March 2010.

[13] Rick Rabinovich, "40 Gb/s and 100 Gb/s Ethernet Short Reach Optical and Copper Host Board Channel Design," *IEEE Communications Magazine*, Volume 50, Issue 4, April 2012.

[14] Timothy Prickett Morgan, "IEEE Gets Behind 25G Ethernet Effort," *Enterprise Tech*, July 27, 2014.

[15] Rick Merritt, "50G Ethernet Debate Brewing," *EE Times*, September 3, 2014.

[16] Tom Nolle, "Will We Ever Need 400 Gigabit Ethernet Enterprise Networks?" *Network Evolution E-Zine*, December 2012.

[17] John D'Ambrosia, Paul Mooney, and Mark Nowell, "400 Gb/s Ethernet: Why Now?" *Ethernet Alliance White Paper*, April 2013.

[18] Stephen Hardy, "400 Gigabit Ethernet Task Force Ready to Get to Work," *Lightwave*, March 28, 2014.

[19] John D'Ambrosia, "400GbE and High Performance Computing," *Scientific Computing Blog*, April 18, 2014.

[20] Jim Duffy, "100-Gigabit Ethernet: Bridge to Terabit Ethernet," *Network World*, April 20, 2009.

[21] John D'Ambrosia, "TEF 2014: The Rate Debate," *Ethernet Alliance Blog*, June 23, 2014.

[22] Scott Kipp, "5 New Speeds – 2.5, 5, 25, 50 and 400 GbE," *Ethernet Alliance Blog*, August 8, 2014.

[23] William Stallings, "Gigabit Wi-Fi," *The Internet Protocol Journal*, Volume 17, No. 1, September 2014.

[24] David Chalupsky and Adam Healey, "Datacenter Ethernet: Know Your Options," *Network Computing*, March 28, 2014.

[25] Rich Seifert, *Gigabit Ethernet: Technology and Applications for High-Speed LANs*, ISBN 0-201-18553-9, Addison-Wesley, 1998. (Reviewed in *The Internet Protocol Journal*, Volume 1, Number 2, September 1998.)

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

[27] William Stallings, "Introduction to 5G Part Two: Core Network, Radio Access Network, and Air Interface," *The Internet Protocol Journal*, Volume 26, No. 3, December 2023.

[28] IEEE Standards Association, "IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems," IEEE 1588-2008, July 24, 2008.

[29] IEEE Standards Association, "IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems," IEEE 1588-2019, June 16, 2020.

[30] Ultra Ethernet Consortium: **`https://ultraethernet.org/`**

[31] Wikipedia article on the IEEE 802.3 working group. Lists all current standards, as well as standards under development: **https://en.wikipedia.org/wiki/IEEE_802.3**

[32] Charles Spurgeon and Joann Zimmerman, *Ethernet: The Definitive Guide: Designing and Managing Local Area Networks 2nd Edition*, ISBN-13 978-1449361846, O'Reilly Media, 2014.

[33] Wikipedia article on Exponential Backoff: **https://en.wikipedia.org/wiki/Exponential_backoff**

MIKAEL HOLMBERG is a Distinguished Engineer and member of the office of the CTO at Extreme Networks. He is an experienced professional in networking architectures and technologies including cloud who has worked in the industry for over 30 years and is active in a variety of industry committees.
E-mail: **mikael@extremenetworks.com**

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.

Letter to the Editor

Craig Partridge's "Why ATM Failed" article in *The Internet Protocol Journal*, Volume 26, No 3, December 2023, is excellent. It prompted me to write down a few of my own recollections at the time:

I had thought that it was the work of Sandy Fraser in Bell Labs in 1974 that transformed the *Time Division Multiplexor* (TDM) theory with a label attached to the data segment that identified the now virtual timeslot of each data stream that was the precursor of ATM. My reading on this topic had noted that when this technology was presented to the Bell telephone operating companies as a scalable switching architecture that had far greater flexibility and efficiency than TDM-based architectures in the mid-1980s, the response was positive, and many expected the migration to ATM in the telephone switching fabric to be completed by 2020!

I also recall some material from Dave Sincowski and Bob Lyon from the late 1980s about the early days of ATM. At that time the experience of changing the local network architecture from common-bus 10-Mbps Ethernet to 100-Mbps *Fiber Distributed Data Interface* (FDDI) rings was challenging: It involved large-scale replacement of both the physical network media of the *Local-Area Network* (LAN) and the network interfaces in the attached workstations. The underlying concern was that the next incremental step in LAN speeds would require a similar comprehensive replacement. They were searching for a scalable LAN architecture that could offer increasing capacity without enforced replacement of large parts of the network infrastructure. Bell Lab's Dave Sincowski proposed ATM to Sun Microsystem's Bob Lyons as an answer to that problem, as Sun Microsystem's workstation products had just gone through the Ethernet-to-FDDI NIC transformation. That proposal appears to be why it was the computing sector, not the telephone sector, that was behind the initial adoption of ATM in digital networking.

It wasn't just the small buffer size in ATM switches that was an issue here for high-speed computer networks, it was the choice of the ATM cell size of 53 bytes that proved to be a problem. The cell size decision was already a compromise between a small cell size that more closely matched the TDM slot size that was ideally suited to low-jitter, low-volume voice data streams and a much larger cell size that reduced the per-cell processing overheads of a higher-speed data stream. Switching equipment initially struggled to achieve switching performance of a 10-Mbps Ethernet with a theoretical maximum packet rate of some 1,440 packets per second, so a smaller cell size would require faster processing capability in the switches.

I also recall at that time extensive debate about the "correct" internal buffer dimensioning for ATM switches. Low-jitter objectives call for very shallow internal buffers, while the congestion-loss algorithms of TCP called for delay-bandwidth–sized internal buffers.

I dimly recall a report on Doug Comer's experiments on achievable throughputs using a 155-Mbps ATM switch where the shallow buffers in ATM switches resulted in an achieved 3-Mbps data throughput from a 155-Mbps switch. The LAN market had already gained extensive experience with Ethernet switching, and ATM simply was not an effective alternative in price and performance for local networks.

Despite these issues, for a while in the early to mid 1990s ATM had some success. I vaguely recall the first 155-Mbps backbone circuits in the US used Fore ATM switches with ATM as the only available technology with a clock that was faster than 45 Mbps.

The telephone companies clung to the dream of a single multipurpose digital switched foundation for a suite of data products. The company I worked for at the time in Australia, Telstra, was using Nortel Magellan Passport ATM switches as the basis for their suite of data products as well as telephone trunks. But it was useful for only a brief period of time. When IP services wanted to use circuit transmission rates in excess of 622 Mbps, the Passport switches could not deliver, and at that point the IP product moved to sit beside the Passports and connect to the *Synchronous Digital Hierarchy* (SDH) network at first, and then directly to optical transponders shortly thereafter.

—*Geoff Huston*, **gih@apnic.net**

Check your Subscription Details!

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

Fragments

IAB Statement on Encryption and Mandatory Client-side Scanning of Content

The *Internet Architecture Board* (IAB) recently issued the following statement: "A secure, resilient, and interoperable Internet benefits the public interest and supports human rights[1] to privacy and freedom of opinion and expression. This is endangered by technologies, such as recent proposals for client-side scanning, that mandate unrestricted access to private content and therefore undermine end-to-end encryption and bear the risk to become a widespread facilitator of surveillance and censorship.

This statement is a reaction to recent policy proposals in the United Kingdom[2], European Union[3], United States[4], and other countries that are mandating client-side scans that require access to otherwise end-to-end encrypted content. These proposals envision client-side scanning technologies that search content on devices before it is encrypted or after decryption on receipt. This would potentially be accomplished by comparison against a database maintained by an authority or by leveraging machine learning to identify previously unseen but potentially prohibited content. These envisioned mechanisms fail to consider their broader implications for Internet security.

The *Internet Engineering Task Force* (IETF) is the leading standards development organization for the global Internet. The IAB provides long-range technical direction for Internet development, ensuring the Internet continues to grow and evolve as a platform for global communication and innovation. To create and maintain the Internet as the bedrock of current secure communication, the IETF and the IAB serve as stewards of the Internet's communication protocols and its core values of trust, openness, and fairness that underpin secure online communication. This is accomplished through a transparent process backed by consensus that is open for anybody to participate in. We encourage the continued deployment and strengthening of mechanisms that enhance privacy and security for all users of the Internet.

The IETF and the IAB have published concerns about standardizing wiretaps[5], backdoors[6,7], and surveillance[8], because these technologies reduce the security of the Internet as a whole, fail to curtail malicious actors, and reduce security for Internet users. To ensure all communication can remain properly protected, the IETF continues to develop and enhance encrypted protocols like *Internet Protocol Security* (IPsec)[9] at the IP layer, *Transport Layer Security* (TLS)[10] at the transport layer which is further incorporated into the *Hypertext Transfer Protocol Version 2* (HTTP/2)[11] and QUIC[12] protocols, and inside many application protocols such as email *Secure/Multipurpose Internet Mail Extensions* (S/MIME)[13], *Open Specification for Pretty Good Privacy* (OpenPGP)[14] or instant messaging *Messaging Layer Security* (MLS)[15] and *End-to-End Signing and Object Encryption for the Extensible Messaging and Presence Protocol* (XMPP)[16]. Recognizing that management of increasingly encrypted networks can pose operational challenges, the IAB has recently held a workshop on techniques for managing encrypted networks in ways that intend not to sacrifice security for the Internet's end-users[17].

The IAB has recognized surveillance of any form as a threat to Internet user privacy, where "surveillance is the observation or monitoring of an individual's communications or activities"[18]. As the IAB and *Internet Engineering Steering Group* (IESG) documented in 1996[6], instituting governmental control into communication "provide[s] only a marginal or illusory benefit to law enforcement agencies" as any seemingly beneficial purpose can be equally used by malevolent actors or future authoritarian shifts in government administrations. The IETF community still holds true to these principles today.

For technologies where the intended purpose is scanning of user communication, there is by design no technical way to limit the scope and intent of scanning, nor curtail subsequent changes in scope or intent. Further, specifically when scanning for illegal content, transparency cannot be provided. Mandating such technologies impacts all users of the global Internet and creates a tool that is straightforward to abuse as a widespread facilitator of surveillance and censorship, presenting real-world dangers to the free flow of information and the security and privacy of people. Without privacy, users cannot benefit from the Internet's virtue to connect people and support freedom of expression.

Additionally, one of the founding principles of the Internet has been its openness; the ability for any standards-compliant software to access the network of networks has been the catalyst for world-changing innovations over many decades. Mandatory use of client-side scanning, and the regulatory burden it would impose, would negatively impact this, restrict use of open-source software, and lead to a stagnant landscape where users lose choice.

The IAB shares concerns about societal harms through the distribution of illegal content and criminal action on the Internet and recognizes the need to protect Internet users from such threats. However, the IAB believes that mandating client-side scanning is in direct opposition to the safe, secure and open communication platform that the Internet provides today and undermines the core principles applied by the IAB and the IETF[5, 6, 18] in order to secure the Internet through encryption. The IAB opposes technologies that foster surveillance as they weaken the user's expectations of private communication which decreases the trust in the Internet as the core communication platform of today's society. Mandatory client-side scanning creates a tool that is straightforward to abuse as a widespread facilitator of surveillance and censorship. Mandating on-device scanning of content will compromise privacy, weaken security, and imperil human rights to communication, freedom of expression and freedom of opinion."

References

[1] United Nations Human Rights Office of The High Commissioner, "A/HRC/29/32: Report on encryption, anonymity, and the human rights framework," May 2015.

[2] UK Parliament, "Online Safety Act 2023."

[3] European Commission, "Proposal for a Regulation of The European Parliament and of The Council laying down rules to prevent and combat child sexual abuse," May, 2022.

[4] United States Congress, S.1207, "Eliminating Abusive and Rampant Neglect of Interactive Technologies Act of 2023."

[5] IAB and IETF, "IETF Policy on Wiretapping," RFC 2804, May 2000. *This document articulates why the IETF stated that it was not appropriate to accommodate wiretapping.*

[6] IAB and IESG, "IAB and IESG Statement on Cryptographic Technology and the Internet," RFC 1984, August 1996. *This document stated the IESG and IAB's position regarding legal constraints on encryption in 1996, with a focus on the effects on the Internet. The publication of the document was prompted in part by the controversy surrounding the US government's promotion of the Clipper Chip.*

[7] Jeffrey I. Schiller, "Strong Security Requirements for Internet Engineering Task Force Standard Protocols," RFC 3365, August 2002. *This document set a requirement for IETF standard protocols to use 'appropriate strong security mechanisms,' including encryption. It was published as Best Current Practice in 2002.*

[8] Stephen Farrell and Hannes Tschofenig, "Pervasive Monitoring Is an Attack," RFC 7258, May 2014. *This RFC documents the IETF consensus that pervasive monitoring is an attack, and thus should be mitigated in IETF protocols (often, using encryption). It was a response to the Snowden revelations and an output of the workshop on Strengthening the Internet Against Pervasive Monitoring (STRINT), held jointly by the W3C and IAB.*

[9] IETF IPSec Working Group

[10] Eric Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.3," RFC 8446, August 2018.

[11] Mike Belshe, Roberto Peon, and Martin Thomson, "Hypertext Transfer Protocol Version 2 (HTTP/2)," RFC 7540, May 2015.

[12] IETF QUIC Working Group

[13] IETF S/MIME Mail Security Working Group

[14] IETF Open Specification for Pretty Good Privacy Working Group

[15] IETF Messaging Layer Security Working Group

[16] Peter Saint-Andre, "End-to-End Signing and Object Encryption for the Extensible Messaging and Presence Protocol (XMPP)," RFC 3923, October 2004.

[17] Mallory Knodel, Wes Hardaker, and Tommy Pauly, "Report from the IAB Workshop on Management Techniques in Encrypted Networks (M-TEN)," RFC 9490, January 2024.

[18] Alissa Cooper, Hannes Tschofenig, Bernard Aboba, Jon Peterson, John B. Morris, Jr., Marit Hansen, and Rhys Smith, "Privacy Considerations for Internet Protocols," RFC 6973, July 2013.

The full IAB statement, including additional references, can be found on the IETF datatracker website.

Achieving Greater Heights for MANRS

The *Global Cyber Alliance* (CGA) recently announced a new phase for *Mutually Agreed Norms for Routing Security* (MANRS)[0,1]. The *Internet Society* has partnered with the GCA an international non-profit specializing in addressing cybersecurity challenges at scale by mobilizing stakeholders toward collective action. As part of this partnership, the GCA will take on the functions of the MANRS secretariat and operations, while the Internet Society will maintain significant funding, advocacy, and training functions over the next five years.

In 2014, the Internet Society recognized the industry's willingness for collaborative agreement on best practices for routing security and helped initial participants to capture and share those practices in what became MANRS. Since then, the Internet Society has advocated globally for MANRS uptake, encouraged industry collaboration, supported the evolution of the norms, and evolved to become the secretariat of MANRS.

Fast forward a decade, and MANRS has grown from nine original operators to a community of more than 1,000 participants ranging from small enterprise networks to Tier-1 transit providers, from *Internet Exchange Points* (IXPs) of various sizes to *Content Delivery Network* (CDN) and cloud providers publicly professing their commitment to the MANRS requirements. MANRS is now globally recognized as a beacon for securing global routing.

As MANRS matured, so did the community-led governance model with the establishment of the community-elected Steering Committee. The Internet Society has proudly served as the secretariat, in addition to supporting the initiative with both financial and staff resources as well as operations support to ensure MANRS' growth. In 2019 the *MANRS Observatory*, a conformance measurement tool for routing security, was launched. Since then, many new features have been added to the MANRS Observatory, such as alerts and monthly MANRS readiness reports. Growth also happened through capacity building, and over the years, thousands of network engineers have gone through online courses, virtual labs, and on-site workshops. In 2020, the Internet Society, together with the MANRS community, launched the *Mentors and Ambassadors* program promoting routing security in the areas of research, policy, and training.

MANRS has more than one thousand participating operators across three programs, as well as six network equipment vendors. The initiative has been a tremendous success, but the task of supporting MANRS has grown well beyond the scope of what was a startup initiative 10 years ago. This partnership is an important evolution of a successful initiative that the Internet Society launched, incubated, and nurtured. GCA is honored and excited to step into this role and provide the basis for the long-term sustainability and evolution of MANRS.

Routing security is one of the focus areas of GCA, and the Internet Society and GCA have been collaborating around MANRS since 2021 with excellent results. GCA conducted a survey of network operators to learn more about the state of routing security implementation, the level of concern within network operations and business decision-making, and plans for next steps.

The Internet Society is dedicated to improving routing security and ensuring the best future for MANRS. Over the next five years, the Internet Society will focus on funding and support through training and global advocacy activities, while GCA will provide the secretariat function and operate the MANRS Observatory. GCA is uniquely placed to lead the next evolution of MANRS as its focus on building communities to collectively drive action towards addressing cybersecurity challenges at scale allows it to step into this role and provide the best future home for the operational growth MANRS is experiencing.

GCA is committed to maintaining the vision of MANRS and continuing to expand its global impact. With this partnership, MANRS will continue to achieve greater heights and be further established as the globally recognized benchmark for global routing security.

The partnership builds on the strengths of both organizations—GCA's global footprint of mobilizing communities towards collective action to deliver a secure, trustworthy Internet that enables social and economic progress for all, and the Internet Society's training and advocacy expertise. Together, the Internet Society and GCA are committed to maintaining and expanding the vision of MANRS to continue to increase the awareness and uptake of MANRS principles and improve the Internet's functional integrity.

Everyone who runs a network has a responsibility to ensure a globally robust and secure routing infrastructure. Your network's safety depends on a routing infrastructure that stops bad actors and mitigates accidental misconfigurations that wreak havoc on the Internet. The more network operators work together, the fewer incidents there will be, and the less damage they can do.

For more information about this partnership visit:

https://www.globalcyberalliance.org/achieving-greater-heights-manrs/

References

[0] MANRS: **https://manrs.org/**

[1] Andrei Robachevsky, "Improving Routing Security," *The Internet Protocol Journal*, Volume 22, No. 2, July 2019.

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
Radu Cristian Gheorghiu
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
Michael Hippert
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
Vladislav Kalinovsky
Naoki Kambe
Akbar Kara
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
Christian de Larrinaga
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
Mario Lopez
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
Gavin McCullagh
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
Joseph Moran
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 Preston
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.

Follow us on X and Facebook **@protocoljournal** **https://www.facebook.com/newipj**

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