NOTE: Chapters 1 and 4 is available

1

ELEMENTS OF MODERN NETWORKING

Unit Structure:

1.1 OBJECTIVES

After studying this chapter, you should be able to:

- Explain the key elements and their relationships of a modern networking ecosystem, including end users, network providers, application providers and application service providers.
- Discuss the motivation for the typical network hierarchy of access networks, distribution networks, and core networks.
- Present an overview of Ethernet, including a discussion of its application areas and common data rates.
- Present an overview of Wi-Fi, including a discussion of its application areas and common data rates.
- Understand the differences between the five generations of cellular networks.
- Present an overview of cloud computing concepts.
- Describe the Internet of Things.
- Explain the concepts of network convergence and unified communications.

1.2 THE NETWORKING ECOSYSTEM

As enterprises persist to adjust to the everchanging working nature, understanding the long-term effects of the any road-blocker and what actions we all need to take is critical. As technologists modernize their IT infrastructure, they face a host of obstacles, including legacy infrastructure, poor system integration and teams whose programming skills are not up to snuff. Whether IT professionals need to deliver new applications or create a more efficient IT environment, outmoded IT gets in the way. By contrast, a modern infrastructure adapts, helping IT pros keep pace with business needs.

To create a modern and responsive infrastructure, IT teams have virtualized datacenter infrastructure, from servers to storage to networking. Combined with analytics, virtualization helps the network stand up to new demands. As hardware becomes virtualized and programmable through software, IT teams need the skills to enable integration – and the hardware needs to be able to integrate. How IT professionals choose their tools and design their IT environments to be automated, virtualized, programmable, secure, and scalable will surely be the difference between success and failure.

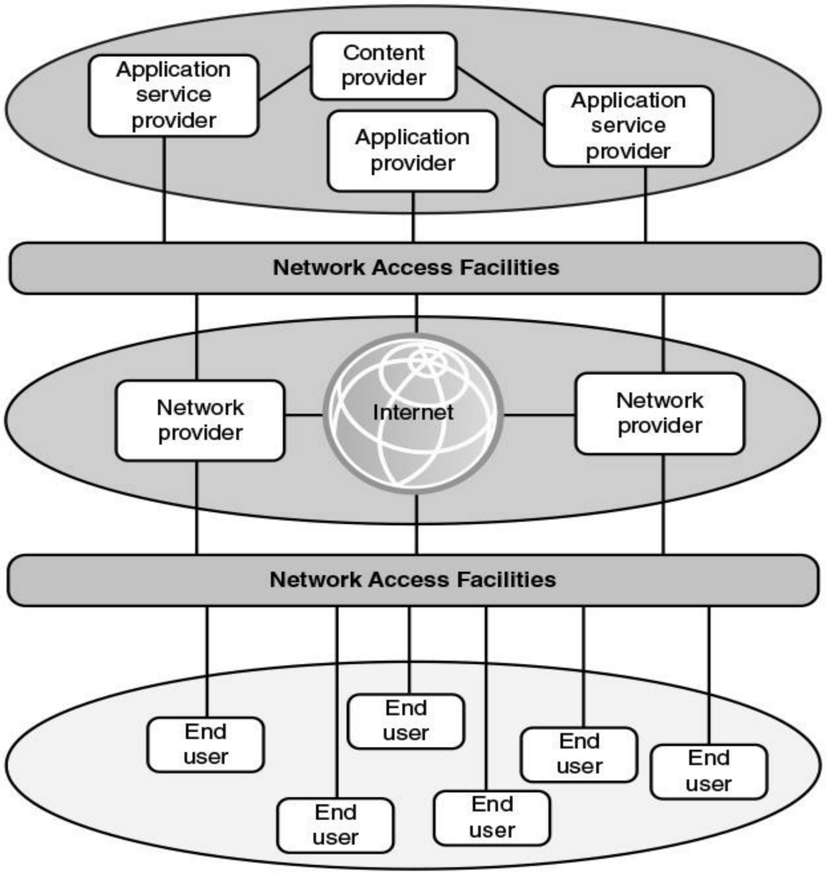

Figure 1.0 The Modern Networking Ecosystem

Figure 1.0 depicts the modern networking ecosystem in very general terms. The entire ecosystem exists to provide services to end users. The term end user, or simply user, is used here as a very general term, to encompass users working within an enterprise or in a public setting or at home. The user platform can be fixed (for example, PC or workstation), portable (for example, laptop), or mobile (for example, tablet or smartphone).

Users connect to network-based services and content through a wide variety of network access facilities. These include digital subscriber line (DSL) and cable modems, Wi-Fi, and Worldwide Interoperability for Microwave Access (WiMAX) wireless modems, and cellular modems. Such network access facilities enable the use to connect directly to the Internet or to a variety of network providers, including Wi-Fi networks, cellular networks, and both private and shared network facilities, such as a premises enterprise network. Ultimately, of course, users want to use network facilities to access applications and content.

Figure 1.0 indicates three broad categories of interest to users. Application providers provide applications, or apps, that run on the user's platform, which is typically a mobile platform. More recently, the concept of an app store has become available for fixed and portable platforms as well.

A distinct category of provider is the application service provider. Whereas the application provider downloads software to the user's platform, the application service provider acts as a server or host of application software that is executed on the provider's platforms. Traditional examples of such software include web servers, e-mail servers, and database servers. The most prominent example now is the cloud computing provider.

The final (topmost) element shown in Figure 1.0 is the content provider. A content provider serves the data to be consumed on the user device (for example, e-mail, music, video). This data may be commercially provided intellectual property. In some instances, an enterprise may be an application or content provider. Examples of content providers are music record labels and movie studios.

Figure 1.0 is intended to provide a very general depiction of the networking ecosystem. It is worth pointing out here two major elements of modern networking not explicitly depicted in this figure:

- **Data center networking:** Both large enterprise data centers and cloud provider data centers consist of very large numbers of interconnected servers. Typically, as much as 80 percent of the data traffic is within the data center network, and only 20 percent relies on external networks to reach users.

  Results of the 2020 Gartner Magic Quadrant (series of market research reports) for Data Center and Cloud Networking is presented below for reference along with an understanding of the criteria for each category:
  - **Leaders** – Cisco, Arista Networks, Juniper Networks (Typically, innovative giants who excel at both vision and execution)
  - **Challengers** – Huawei (Strong execution but low vision)
  - **Visionaries** – VMWare, Dell EMC, Cumulus Networks, HPE (Aruba) (Good vision but low execution)
  - **Niche Players** – NVIDIA-Mellanox Technologies, Extreme, H3C (Hyper-focused on a small segment, resulting in low vision and low execution)

- **IoT or fog networking:** An Internet of Things deployed by an enterprise may consist of hundreds, thousands, even millions of devices. The vast bulk of the data traffic to and from these devices is machine to machine, rather than user to machine.

Each of these networking environments creates its own requirements, which are discussed as the book progresses.

1.3 EXAMPLE NETWORK ARCHITECTURES

This section introduces two example network architectures, and with them some of the networking terminology in common use.

1.3.1 A Global Network Architecture

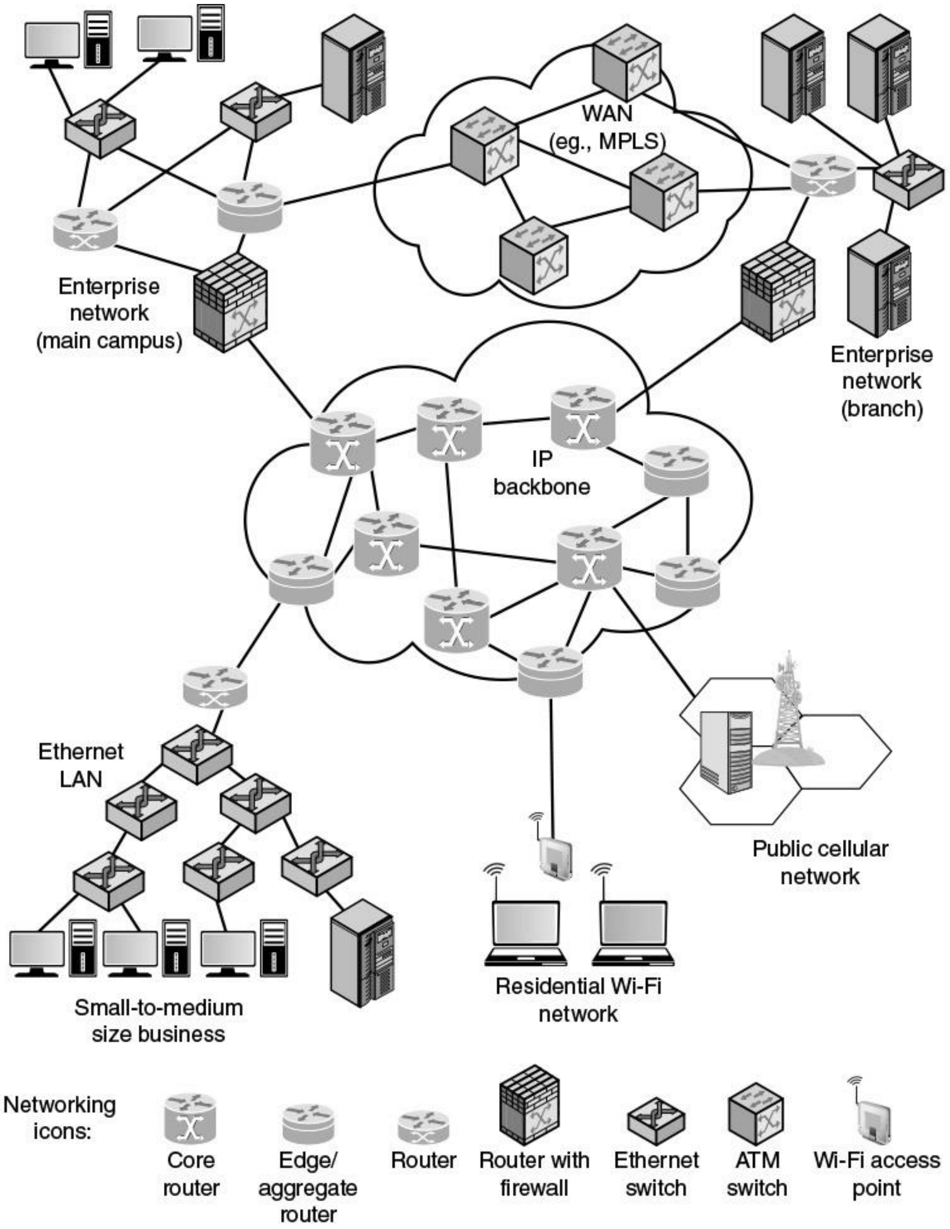

Figure 1.1 A Global Networking Architecture

We begin with an architecture that could represent an enterprise network of national or global extent, or a portion of the Internet with some of its associated networks. Figure 1.1 illustrates some of the typical communications and network elements in use in such a context.

At the center of the figure is an IP backbone, or core, network, which could represent a portion of the Internet or an enterprise IP network. Typically, the backbone consists of high-performance routers, called core routers, interconnected with high-volume optical links. The optical links often make use of what is known as wavelength-division multiplexing (WDM), such that each link has multiple logical channels occupying different portions of the optical bandwidth.

At the periphery of an IP backbone are routers that provide connectivity to external networks and users. These routers are sometimes referred to as edge routers or aggregation routers. Aggregation routers are also used within an enterprise network to connect several routers and switches, to external resources, such as an IP backbone or a high-speed WAN. As an indication of the capacity requirements for core and aggregation routers, the IEEE Ethernet Bandwidth Assessments Group [XI11] reports on an analysis that projects these requirements for Internet backbone providers and large enterprise networks in China. The analysis concludes that aggregation router requirements will be in the range of 200 Gbps to 400 Gbps per optical link by 2020, and 400 Gbps to 1 Tbps per optical link for core routers by 2020.

The upper part of Figure 1.1 depicts a portion of what might be a large enterprise network. The figure shows two sections of the network connected via a private high-speed WAN, with switches interconnected with optical links. MPLS using IP is a common switching protocol used for such WANs; wide-area Ethernet is another option. Enterprise assets are connected to, and protected from, an IP backbone or the Internet via routers with firewall capability, a not uncommon arrangement for implementing the firewall. The lower left of the figure depicts what might be a layout for a small- or medium-size business, which relies on an Ethernet LAN. Connection to the Internet through a router could be through a cable or DSL connection or a dedicated high-speed link.

The lower portion of Figure 1.1 also shows an individual residential user connected to an Internet service provider (ISP) through some sort of subscriber connection. Common examples of such a connection are a DSL, which provides a high-speed link over telephone lines and requires a special DSL modem, and a cable TV facility, which requires a cable modem, or some type of wireless connection. In each case, there are separate issues concerning signal encoding, error control, and the internal structure of the subscriber network. Finally, mobile devices, such as smartphones and tablets, can connect to the Internet through the public cellular network, which has a high-speed connection, typically optical, to the Internet.

1.3.2 A Typical Network Hierarchy

This section focuses in on a network architecture that, with some variation, is common in many enterprises. As Figure 1.2 illustrates, enterprises often design their network facilities in a three-tier hierarchy: access, distribution, and core.

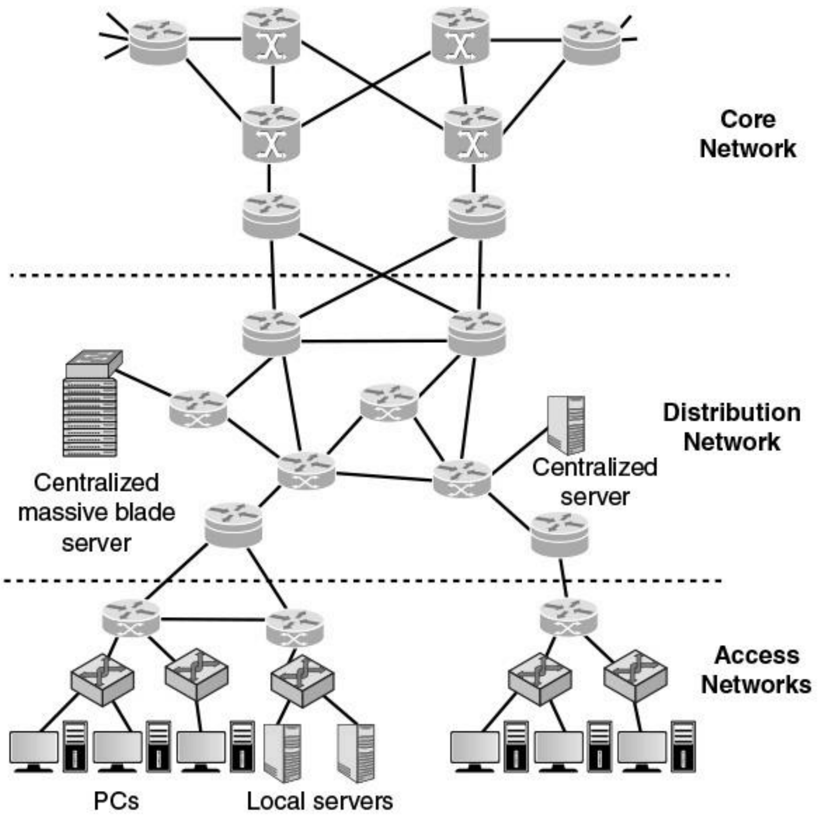

Figure 1.2 A Typical Network Hierarchy

Closest to the end user is the access network. Typically, an access network is a local-area network (LAN) or campus-wide network that consisting of LAN switches (typically Ethernet switches) and, in larger LANs, IP routers that provide connectivity among the switches. Layer 3 switches (not shown) are also commonly used within an LAN. The access network supports end user equipment, such as desktop and laptop computers and mobile devices. The access network also supports local servers that primarily or exclusively serve the users on the local access network.

One or more access routers connect the local assets to the next higher level of the hierarchy, the distribution network. This connection may be via the Internet or some other public or private communications facility. Thus, as described in the preceding subsection, these access routers function as edge routers that forward traffic into and out of the access network. For a large local facility, there might be additional access routers that provide internal routing but do not function as edge routers (not shown in Figure 1.1).

The distribution network connects access networks with each other and with the core network. An edge router in the distribution network connects to an edge router in an access network to provide connectivity. The two routers are configured to recognize each other and will generally exchange routing and connectivity information and, typically, some traffic-related information. This cooperation between routers is referred to as peering. The distribution network also serves to aggregate traffic destined for the core router, which protects the core from high-density peering. That is, the use of a distribution network limits the number of routers that establish peer relationships with edge routers in the core, saving memory, processing, and transmission capacity. A distribution network may also directly connect servers that are of use to multiple access networks, such as database servers and network management servers.

Again, as with access networks, some of the distribution routers may be purely internal and do not provide an edge router function. The core network, also referred to as a backbone network, connects geographically dispersed distribution networks as well as providing access to other networks that are not part of the enterprise network. Typically, the core network will use very high-performance routers, high-capacity transmission lines, and multiple interconnected routers for increased redundancy and capacity. The core network may also connect to high-performance, high-capacity servers, such as large database servers and private cloud facilities.

Some of the core routers may be purely internal, providing redundancy and additional capacity without serving as edge routers.

A hierarchical network architecture is an example of a good modular design. With this design, the capacity, features, and functionality of network equipment (routers, switches, network management servers) can be optimized for their position in the hierarchy and the requirements at a given hierarchical level.

1.4. ETHERNET

The concept of Ethernet was formulated and introduced by XEROX PARC, now simply known as PARC (Palo Alto Research Centre).This agency proposed to develop a form of system that would permit/allow computers and devices to be connected with one and other using coaxial cables. Engineers Bob Metcalfe and D.R Boggs developed Ethernet beginning in 1972. In 1976, a connection two computers were made, and data transfer fruitfully took place with the speed of 3MB/second. In 1980, industry standards based on their work were established under IEEE 802.3 set of specifications. In 1990’s, fast Ethernet technology came into existence fulfilling the objective of:

a) increasing the performance of previous traditional Ethernet
b) avoiding the need of completely re-cable existing Ethernet networks.

Technologies like Ethernet, Wi-Fi, and 4G/5G cellular networks have evolved to support very high data rates supporting many multimedia applications required by enterprises, consumers and, at the same time, place great demands on network switching equipment and network management facilities.

1.4.1 Applications of Ethernet

Ethernet is the predominant wired networking technology, used in homes, offices, data centers, enterprises, and WANs. As Ethernet has evolved to support data rates up to 100 Gbps and distances from a few meters to tens of kilometers, it has become essential for supporting personal computers, workstations, servers, and massive data storage devices in organizations large and small.

Ethernet in the Home

Ethernet has long been used in the home to create a local network of computers with access to the Internet via a broadband modem/router. With the increasing availability of high-speed, low-cost Wi-Fi on computers, tablets, smartphones, modem/routers, and other devices, home reliance on Ethernet has declined. Nevertheless, almost all home networking setups include some use of Ethernet.

Two recent extensions of Ethernet technology have enhanced and broadened the use of Ethernet in the home: powerline carrier (PLC) and Power over Ethernet (PoE). Powerline modems take advantage of existing power lines and use the power wire as a communication channel to transmit Ethernet packets on top of the power signal. This makes it easy to include Ethernet-capable devices throughout the home into the Ethernet network.

PoE acts in a complementary fashion, distributing power over the Ethernet data cable. PoE uses the existing Ethernet cables to distribute power to devices on the network, thus simplifying the wiring for devices such as computers and televisions. With all these Ethernet options, Ethernet will retain a strong presence in home networking, complementing the advantages of Wi-Fi.

Ethernet in the Office

Ethernet has also long been the dominant network technology for wired local-area networks (LANs) in the office environment. Early on there were some competitors, such as IBM’s Token Ring LAN and the Fiber Distributed Data Interface (FDDI), but the simplicity, performance, and wide availability of Ethernet hardware eventually made Ethernet the winner. Today, as with home networks, the wired Ethernet technology exists side by side with the wireless Wi-Fi technology. Much of the traffic in a typical office environment now travels on Wi-Fi, particularly to support mobile devices. Ethernet retains its popularity because it can support many devices at high speeds, is not subject to interference, and provides a security advantage because it is resistant to eavesdropping. Therefore, a combination of Ethernet and Wi-Fi is the most common architecture.

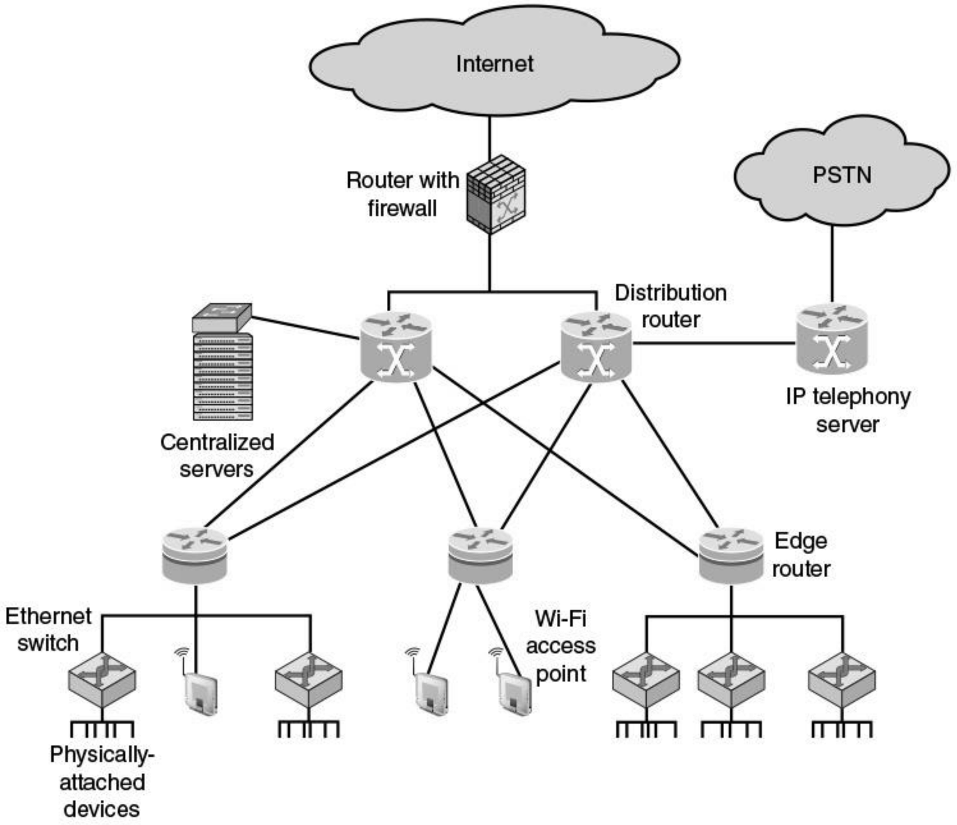

Figure 1.3 A Basic Enterprise LAN architecture

Figure 1.3 provides a simplified example of an enterprise LAN architecture. The LAN connects to the Internet/WANs via a firewall. A hierarchical arrangement of routers and switches provides the interconnection of servers, fixed user devices, and wireless devices. Typically, wireless devices are only attached at the edge or bottom of the hierarchical architecture; the rest of the campus infrastructure is all Ethernet. There may also be an IP telephony server that provides call control functions (voice switching) for the telephony operations in an enterprise network, with connectivity to the public switched telephone network (PTSN).

Ethernet in the Enterprise

A tremendous advantage of Ethernet is that it is possible to scale the network, both in terms of distance and data rate, with the same Ethernet protocol and associated quality of service (QoS) and security standards. An enterprise can easily extend an Ethernet network among several buildings on the same campus or even some distance apart, with links ranging from 10 Mbps to 100 Gbps, using a mixture of cable types and Ethernet hardware. Because all the hardware and communications software conform to the same standard, it is easy to mix different speeds and different vendor equipment. The same protocol is used for intensive high-speed interconnections of data servers in a single room, workstations and servers distributed throughout the building, and links to Ethernet networks in other buildings up to 100 km away.

Ethernet in the Data Center

As in other areas, Ethernet has come to dominate in the data center, where very high data rates are needed to handle massive volumes of data among networked servers and storage units. Historically, data centers have employed various technologies to support high-volume, short-distance needs, including InfiniBand and Fiber Channel. But now that Ethernet can scale up to 100 Gbps, with 400 Gbps on the horizon, the case for a unified protocol approach throughout the enterprise is compelling. Two features of the new Ethernet approach are noteworthy. For co-located servers and storage units, high-speed Ethernet fiber links and switches provided the needed networking infrastructure. Another important version of Ethernet is known as backplane Ethernet. Backplane Ethernet runs over copper jumper cables that can provide up to 100 Gbps over very short distances. This technology is ideal for blade servers, in which multiple server modules are housed in a single chassis.

Ethernet for Wide-Area Networking

Until recently, Ethernet was not a significant factor in wide-area networking. But gradually, more telecommunications and network providers have switched to Ethernet from alternative schemes to support wide-area access (also referred to as first mile or last mile). Ethernet is supplanting a variety of other wide-area options, such as dedicated T1 lines, synchronous digital hierarchy (SDH) lines, and Asynchronous Transfer Mode (ATM). When used in this fashion, the term carrier Ethernet is applied. The term metro Ethernet, or metropolitan-area network (MAN) Ethernet, is also used. Ethernet has the advantage that it seamlessly fits into the enterprise network for which it provides wide-area access. But a more important advantage is that carrier Ethernet provides much more flexibility in terms of the data rate capacity that is used, compared to traditional wide-area alternatives. Carrier Ethernet is one of the fastest-growing Ethernet technologies, destined to become the dominant means by which enterprises access wide-area networking and Internet facilities.

1.4.2 Standards

Within the IEEE 802 LAN standards committee, the 802.3 group is responsible for issuing standards for LANs that are referred to commercially as Ethernet. Complementary to the efforts of the 802.3 committee, the industry consortium known as The Ethernet Alliance supports and originates activities that span from incubation of new Ethernet technologies to interoperability testing to demonstrations to education.

1.4.3 Ethernet Data Rates

Currently, Ethernet systems are available at speeds up to 100 Gbps. Here is a brief chronology:

- 1983: 10 Mbps (megabit per second, million bits per second)
- 1995: 100 Mbps
- 1998: 1 Gbps (gigabits per second, billion bits per second)
- 2003: 10 Gbps
- 2010: 40 Gbps and 100 Gbps
- 2017: 700 Gbps

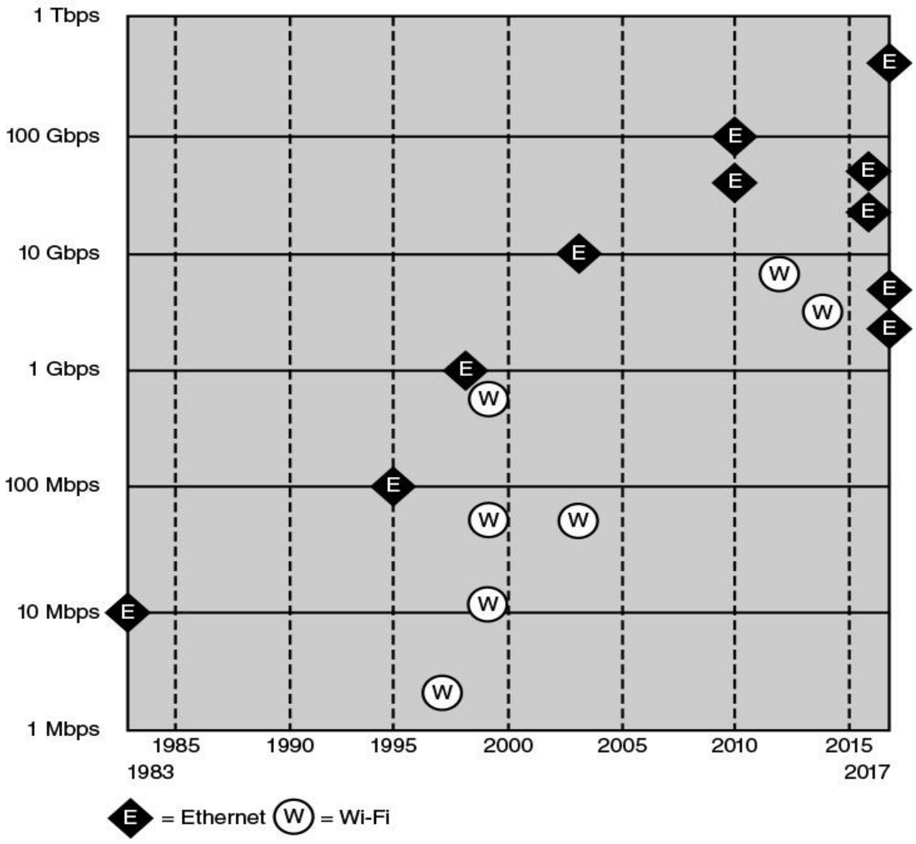

Figure 1.4 Ethernet and Wi-Fi Timelines

1-Gbps Ethernet

For several years, the initial standard of Ethernet, at 10 Mbps, was adequate for most office environments. By the early 1990s, it was clear that higher data rates were needed to support the growing traffic load on the typical LAN. Key drivers included the following:

- Centralized server farms: In many multimedia applications, there is a need for client system to be able to draw huge amounts of data from multiple, centralized servers, called server farms. As the performance of the servers has increased, the network becomes the bottleneck.

- Power workgroups: These groups typically consist of a small number of cooperating users who need to exchange massive data files across the network. Example applications are software development and computer-aided design.

- High-speed local backbone: As processing demand grows, enterprises develop an architecture of multiple LANs interconnected with a high-speed backbone network.

To meet such needs, the IEEE 802.3 committee developed a set of specifications for Ethernet at 100 Mbps, followed a few years later by a 1-Gbps family of standards. In each case, the new specifications defined transmission media and transmission encoding schemes built on the basic Ethernet framework, making the transition easier than if a completely new specification were issued.

10-Gbps Ethernet

Even as the ink was drying on the 1-Gbps specification, the continuing increase in local traffic made this specification inadequate for needs in the short-term future. Accordingly, the IEEE 802.3 committee soon issued a standard for 10-Gbps Ethernet. The principle driving requirement for 10-Gbps Ethernet was the increase in intranet (local interconnected networks) and Internet traffic.

Several factors contribute to the explosive growth in both Internet and intranet traffic:

- An increase in the number of network connections
- An increase in the connection speed of each end-station (for example, 10-Mbps users moving to 100 Mbps, analog 56k users moving to DSL and cable modems).
- An increase in the deployment of bandwidth-intensive applications such as high-quality video.
- An increase in web hosting and application hosting traffic.

Initially, network managers used 10-Gbps Ethernet to provide high-speed, local backbone interconnection between large-capacity switches. As the demand for bandwidth increased, 10-Gbps Ethernet began to be deployed throughout the entire network, to include server farm, backbone, and campus-wide connectivity. This technology enables ISPs and network service providers (NSPs) to create very high-speed links at a very low cost between co-located carrier-class switches and routers.

The technology also allows the construction of MANs and WANs that connect geographically dispersed LANs between campuses or points of presence (PoPs).

100-Gbps Ethernet

The IEEE 802.3 committee soon realized the need for a greater data rate capacity than 10-Gbps Ethernet offers, to support Internet exchanges, high-performance computing, and video-on-demand delivery. The authorization request justified the need for two different data rates in the new standard (40 Gbps and 100 Gbps) by recognizing that aggregate network requirements and end-station requirements are increasing at different rates.

The following are market drivers for 100-Gbps Ethernet:

- **Data center/Internet media providers:** To support the growth of Internet multimedia content and web applications, content providers have been expanding data centers, pushing 10-Gbps Ethernet to its limits. Likely to be high-volume early adopters of 100-Gbps Ethernet.

- **Metro video/service providers:** Video on demand has been driving a new generation of 10-Gbps Ethernet metropolitan/core network buildouts. Likely to be high-volume adopters in the medium term.

- **Enterprise LANs:** Continuing growth in convergence of voice/video/data and in unified communications is driving up network switch demands. However, most enterprises still rely on 1-Gbps or a mix of 1-Gbps and 10-Gbps Ethernet, and adoption of 100-Gbps Ethernet is likely to be slow.

- **Internet exchanges/ISP core routing:** With the massive amount of traffic flowing through these nodes, these installations are likely to be early adopters of 100-Gbps Ethernet.

Figure 1.5 shows an example of the application of 100-Gbps Ethernet. The trend at large data centers, with substantial banks of blade servers, is the deployment of 10-Gbps ports on individual servers to handle the massive multimedia traffic provided by these servers.

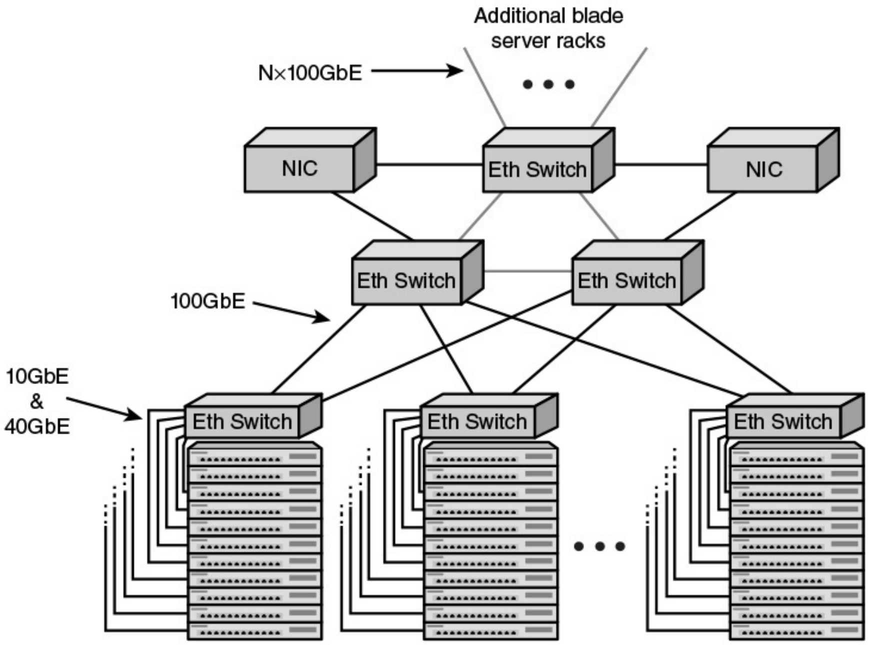

Figure 1.5 Configuration of massive blade server cloud site

Typically, a single blade server rack will contain multiple servers and one or two 10-Gbps Ethernet switches to interconnect all the servers and provide connectivity to the rest of the facility. The switches are often mounted in the rack and referred to as top-of-rack (ToR) switches. The term ToR has become synonymous with server access switch, even if it is not located "top of rack." For very large data centers, such as cloud providers, the interconnection of multiple blade server racks with additional 10-Gbps switches is increasingly inadequate.

To handle the increased traffic load, switches operating at greater than 10 Gbps are needed to support the interconnection of server racks and to provide adequate capacity for connecting offsite through network interface controllers (NICs).

25/50-Gbps Ethernet

One of the options for implementing 100-Gbps is as four 25-Gbps physical lanes. Therefore, it would be relatively easy to develop standards for 25-Gbps and 50-Gbps Ethernet, using one or two lanes, respectively. Having these two lower-speed alternatives, based on the 100-Gbps technology, would give users more flexibility in meeting existing and near-term demands with a solution that would scale easily to higher data rates. Such considerations have led to the form of the 25 Gigabit Ethernet Consortium by several leading cloud networking providers, including Google and Microsoft. The objective of the Consortium is to support an industry-standard, interoperable Ethernet specification that boosts the performance and slashes the interconnect cost per Gbps between the NIC and ToR switch. The specification adopted by the Consortium prescribes a single-lane 25-Gbps Ethernet and dual-lane 50-Gbps Ethernet link protocol, enabling up to 2.5 times higher performance per physical lane on twinax copper wire between the rack endpoint and switch compared to 10-Gbps and 40-Gbps Ethernet links. The IEEE 802.3 committee is at work developing the needed standards for 25 Gbps and may include 50 Gbps.

It is too early to say how these various options (25, 40, 50, 100 Gbps) will play out in the marketplace. In the intermediate term, the 100-Gbps switch is likely to predominate at large sites, but the availability of these slower and cheaper alternatives gives enterprises several paths for scaling up to meet increasing demand.

400-Gbps Ethernet

The growth in demand never lets up. IEEE 802.3 is currently exploring technology options for producing a 400-Gbps Ethernet standard (proposed as 802.3db), although no timetable is yet in place. Looking beyond that milestone, there is widespread acknowledgment that a 1-Tbps (terabits per second, trillion bits per second) standard will eventually be produced.

2.5/5-Gbps Ethernet

As a testament to the versatility and ubiquity of Ethernet, and while ever higher data rates are being standardized, consensus is developing to standardize two lower rates: 2.5 Gbps and 5 Gbps. These relatively low speeds are also known as Multirate Gigabit BASE-T (MGBASE-T). Currently, the MGBASE-T Alliance is overseeing the development of these standards outside of IEEE. It is likely that the IEEE 802.3 committee will ultimately issue standards based on these industry efforts.

These new data rates are mainly intended to support IEEE 802.11ac wireless traffic into a wired network. IEEE 802.11ac is a 3.2-Gbps Wi-Fi standard that is gaining acceptance where more than 1 Gbps of throughput is needed, such as to support mobile users in the office environment. This new wireless standard overruns 1-Gbps Ethernet link support but may not require the next step up, which is 10 Gbps. If 2.5 and 5 Gbps can be made to work over the same cable that supports 1 Gbps, this would provide a much-needed uplink speed improvement for access points supporting 802.11ac radios with their high bandwidth capabilities.

1.5 Wi-Fi

Just as Ethernet has become the dominant technology for wired LANs, so Wi-Fi, standardized by the IEEE 802.11 committee, has become the dominant technology for wireless LANs. This overview section discusses applications of Wi-Fi and then looks at standards and performance.

1.5.1 Applications of Wi-Fi

Wi-Fi is the predominant wireless Internet access technology, used in homes, offices, and public spaces. Wi-Fi in the home now connects computers, tablets, smartphones, and a host of electronic devices, such as video cameras, TVs, and thermostats. Wi-Fi in the enterprise has become an essential means of enhancing worker productivity and network effectiveness. And public Wi-Fi hotspots have expanded dramatically to provide free Internet access in must public places.

Wi-Fi in the Home

The first important use of Wi-Fi in the home was to replace Ethernet cabling for connecting desktop and laptop computers with each other and with the Internet. A typical layout is a desktop computer with an attached router/modem that provides an interface to the Internet. Other desktop and laptop computers connect either via Ethernet or Wi-Fi to the central router, so that all the home computers can communicate with each other and with the Internet. Wi-Fi greatly simplified the hookup. Not only is there no need for a physical cable hookup, but the laptops can be moved easily from room to room or even outside the house.

Today, the importance of Wi-Fi in the home has expanded tremendously. Wi-Fi remains the default scheme for interconnecting a home computer network. Because both Wi-Fi and cellular capability are now standard on both smartphones and tablets, the home Wi-Fi provides a cost-effective way to the Internet. The smartphone or tablet will automatically use a Wi-Fi connection to the Internet if available, and only switch to the more expensive cellular connection if the Wi-Fi connection is not available. And Wi-Fi is essential to implementing the latest evolution of the Internet: Internet of Things.

Public Wi-Fi

Access to the Internet via Wi-Fi has expanded dramatically in recent years, as more and more facilities provide a Wi-Fi hotspot, which enables any Wi-Fi device to attach. Wi-Fi hotspots are provided in coffee shops, restaurants, train stations, airports, libraries, hotels, hospitals, department stores, RV parks, and many other places. So many hotspots are available that it is rare to be too far from one. There are now numerous tablet and smartphone apps that increase their convenience.

Even very remote places will be able to support hotspots with the development of the satellite Wi-Fi hotspot. The first company to develop such a product is the satellite communications company Iridium. The satellite modem will initially provide a relatively low-speed connection, but the data rates will inevitably increase.

Enterprise Wi-Fi

The economic benefit of Wi-Fi is most clearly seen in the enterprise. Wi-Fi connections to the enterprise network have been offered by many organizations of all sizes, including public and private sector. But in recent years, the use of Wi-Fi has expanded dramatically, to the point that now approximately half of all enterprise network traffic is via Wi-Fi rather than the traditional Ethernet. Two trends have driven the transition to a Wi-Fi-centered enterprise. First, the demand has increased, with more and more employees preferring to use laptops, tablets, and smartphones to connect to the enterprise network, rather than a desktop computer. Second, the arrival of Gigabit Ethernet, especially the IEEE 802.ac standard, allows the enterprise network to support high-speed connections to many mobile devices simultaneously.

Whereas Wi-Fi once merely provided an accessory network designed to cover meetings and public areas, enterprise Wi-Fi deployment now generally provides ubiquitous coverage, to include main offices and remote facilities, and both indoor locations and outdoor spaces surrounding them. Enterprises accepted the need for, and then began to encourage, the practice known as bring your own device (BYOD). The almost universal availability of Wi-Fi capability on laptops, tablets, and smartphones, in addition to the wide availability of home and public Wi-Fi networks, has greatly benefited the organization. Employees can use the same devices and the same applications to continue their work or check their e-mail from wherever they are—home, at their local coffee shop, or while traveling. From the enterprise perspective, this means higher productivity and efficiency and lower costs.

1.5.2 Standards

Essential to the success of Wi-Fi is interoperability. Wi-Fi-enabled devices must be able to communicate with Wi-Fi access points, such as the home router, the enterprise access point, and public hotspots, regardless of the manufacturer of the device or access point. Such interoperability is guaranteed by two organizations. First, the IEEE 802.11 wireless LAN committee develops the protocol and signaling standards for Wi-Fi. Then, the Wi-Fi Alliance creates test suites to certify interoperability for commercial products that conform to various IEEE 802.11 standards. The term Wi-Fi (wireless fidelity) is used for products certified by the Alliance.

1.5.3 Wi-Fi Data Rates

Just as businesses and home users have generated a need to extend the Ethernet standard to speeds in the gigabits per second (Gbps) range, the same requirement exists for Wi-Fi. As the technology of antennas, wireless transmission techniques, and wireless protocol design has evolved, the IEEE 802.11 committee has been able to introduce standards for new versions of Wi-Fi at ever-higher speeds. Once the standard is issued, industry quickly develops the products. Here is a brief chronology, starting with the original standard, which was simply called IEEE 802.11, and showing the maximum data rate for each version (Figure 1.4):

- 802.11 (1997): 2 Mbps (megabits per second, million bits per second)
- 802.11a (1999): 54 Mbps
- 802.11b (1999): 11 Mbps
- 802.11n (1999): 600 Mbps
- 802.11g (2003): 54 Mbps
- 802.11ad (2012): 6.76 Gbps (billion bits per second)
- 802.11ac (2014): 3.2 Gbps

IEEE 802.11ac operates in the 5-GHz band, as does the older and slower standards 802.11a and 802.11n. It is designed to provide a smooth evolution from 802.11n. This new standard makes use of advanced technologies in antenna design and signal processing to achieve much greater data rates, at lower battery consumption, all within the same frequency band as the older versions of Wi-Fi.

IEEE 802.11ad is a version of 802.11 operating in the 60- GHz frequency band. This band offers the potential for much wider channel bandwidth than the 5-GHz band, enabling high data rates with relatively simple signal encoding and antenna characteristics. Few devices operate in the 60-GHz band, which means communication experiences less interference than in the other bands used for Wi-Fi.

Because of the inherent transmission limitations of the 60-GHz band, 802.11ad is likely to be useful only within a single room. Because it can support high data rates and, for example, could easily transmit uncompressed high-definition video, it is suitable for applications such as replacing wires in a home entertainment system, or streaming high-definition movies from your cell phone to your television.

Gigabit Wi-Fi holds attractions for both office and residential environments and commercial products are beginning to roll out. In the office environment, the demand for ever greater data rates has led to Ethernet offerings at 10 Gbps, 40 Gbps, and most recently 100 Gbps. These stupendous capacities are needed to support blade servers, heavy reliance on video and multimedia, and multiple broadband connections offsite. At the same time, the use of wireless LANs has grown dramatically in the office setting to meet needs for mobility and flexibility. With the gigabit-range data rates available on the fixed portion of the office LAN, gigabit Wi-Fi is needed to enable mobile users to effectively use the office resources. IEEE 802.11ac is likely to be the preferred gigabit Wi-Fi option for this environment.

In the consumer and residential market, IEEE 802.11ad is likely to be popular as a low-power, short-distance wireless LAN capability with little likelihood of interfering with other devices. IEEE 802.11ad is also an attractive option in professional media production environments in which massive amounts of data need to be moved short distances.

1.6 4G/5G CELLULAR

Cellular technology is the foundation of mobile wireless communications and supports users in locations that are not easily served by wired networks. Cellular technology is the underlying technology for mobile telephones, personal communications systems, wireless Internet, and wireless web applications, and much more. This section looks at how cellular technology has evolved through four generations and is poised for a fifth generation.

1.6.1 First Generation

The original cellular networks, now dubbed 1G, provided analog traffic channels and were designed to be an extension of the public switched telephone networks. Users with brick-sized cell phones placed and received calls in the same fashion as landline subscribers. The most widely deployed 1G system was the Advanced Mobile Phone Service (AMPS), developed by AT&T. Voice transmission was purely analog and control signals were sent over a 10-kbps analog channel.

1.6.2 Second Generation

First-generation cellular networks quickly became highly popular, threatening to swamp available capacity. Second-generation (2G) systems were developed to provide higher-quality signals, higher data rates for support of digital services, and greater capacity. Key differences between 1G and 2G networks include the following:

Digital traffic channels: The most notable difference between the two generations is that 1G systems are almost purely analog, whereas 2G systems are digital. 1G systems are designed to support voice channels; digital traffic is supported only using a modem that converts the digital data into analog form. 2G systems provide digital traffic channels. These systems readily support digital data; voice traffic is first encoded in digital form before transmitting.

- Encryption: Because all the user traffic, and the control traffic, is digitized in 2G systems, it is a relatively simple matter to encrypt all the traffic to prevent eavesdropping. All 2G systems provide this capability, whereas 1G systems send user traffic in the clear, providing no security.

- Error detection and correction: The digital traffic stream of 2G systems also lends itself to the use of error detection and correction techniques. The result can be very clear voice reception.

- Channel access: In 1G systems, each cell supports several channels. At any given time, a channel is allocated to only one user. 2G systems also provide multiple channels per cell, but each channel is dynamically shared by several users.

1.6.3 Third Generation

The objective of the third generation (3G) of wireless communication is to provide high-speed wireless communications to support multimedia, data, and video in addition to voice. 3G systems share the following design features:

- **Bandwidth**: An important design goal for all 3G systems is to limit channel usage to 5 MHz There are several reasons for this goal. On the one hand, a bandwidth of 5 MHz or more improves the receiver's ability to resolve multipath when compared to narrower bandwidths. On the other hand, the available spectrum is limited by competing needs, and 5 MHz is a reasonable upper limit on what can be allocated for 3G. Finally, 5 MHz is adequate for supporting data rates of 144 and 384 kbps, the main targets for 3G services.

- **Data rate**: Target data rates are 144 and 384 kbps. Some 3G systems also provide support up to 2 Mbps for office use.

- **Multirate**: The term multirate refers to the provision of multiple fixed-data-rate logical channels to a given user, in which different data rates are provided on different logical channels. Further, the traffic on each logical channel can be switched independently through the wireless and fixed networks to different destinations. The advantage of multirate is that the system can flexibly support multiple simultaneous applications from a given user and can efficiently use available capacity by only providing the capacity required for each service.

1.6.4 Fourth Generation

The evolution of smartphones and cellular networks has ushered in a new generation of capabilities and standards, which is collectively called 4G. 4G systems provide ultra-broadband Internet access for a variety of mobile devices including laptops, smartphones, and tablets. 4G networks support Mobile web access and high-bandwidth applications such as high-definition mobile TV, mobile video conferencing, and gaming services.

These requirements have led to the development of a fourth generation (4G) of mobile wireless technology that is designed to maximize bandwidth and throughput while also maximizing spectral efficiency. 4G systems have the following characteristics:

- Based on an all-IP packet switched network
- Support peak data rates of up to approximately 100 Mbps for high-mobility mobile access and up to approximately 1 Gbps for low-mobility access such as local wireless access
- Dynamically share and use the network resources to support more simultaneous users per cell
- Support smooth handovers across heterogeneous networks
- Support high QoS for next-generation multimedia applications

In contrast to earlier generations, 4G systems do not support traditional circuit-switched telephone service, providing only IP telephony services.

1.6.5 Fifth Generation

In telecommunications, 5G is the fifth-generation technology standard for broadband cellular networks, which cellular phone companies began deploying worldwide in 2019, and is the planned successor to the 4G networks which provide connectivity to most current cellphones. Like its predecessors, 5G networks are cellular networks, in which the service area is divided into small geographical areas called cells. All 5G wireless devices in a cell are connected to the Internet and telephone network by radio waves through a local antenna in the cell. The main advantage of the new networks is that they will have greater bandwidth, giving higher download speeds, eventually up to 10 gigabits per second (Gbit/s). Due to the increased bandwidth, it is expected that the new networks will not just serve cellphones like existing cellular networks, but also be used as general internet service providers for laptops and desktop computers, competing with existing ISPs such as cable internet, and also will make possible new applications in internet of things (IoT) and machine to machine areas. Current 4G cellphones will not be able to use the new networks, which will require new 5G enabled wireless devices.

The increased speed is achieved partly by using higher-frequency radio waves than current cellular networks. However, higher-frequency radio waves have a shorter range than the frequencies used by previous cell phone towers, requiring smaller cells. So, to ensure wide service, 5G networks operate on up to three frequency bands, low, medium, and high. A 5G network will be composed of networks of up to 3 different types of cells, each requiring different antennas, each type giving a different trade-off of download speed vs. distance and service area. 5G cellphones and wireless devices will connect to the network through the highest speed antenna within range at their location:

- Low band 5G uses a similar frequency range to current 4G cellphones, 600-700 MHz, giving download speeds a little higher than 4G: 30-250 megabits per second (Mbit/s). Low-band cell towers will have a range and coverage area like current 4G towers. Mid-band 5G uses microwaves of 2.5-3.7 GHz, currently allowing speeds of 100-900 Mbit/s, with each cell tower providing service up to several miles in radius. This level of service is the most widely deployed and should be available in most metropolitan areas in 2020. Some countries are not implementing low band, making this the minimum service level.

- High band 5G currently uses frequencies of 25-39 GHz, near the bottom of the millimeter wave band, although higher frequencies may be used in the future. It often achieves download speeds of a gigabit per second (Gbit/s), comparable to cable internet. However, millimeter waves (mmWave or mmW) have a more limited range, requiring many small cells. They have trouble passing through some types of walls and windows. Due to their higher costs, current plans are to deploy these cells only in dense urban environments and areas where crowds of people congregate such as sports stadiums and convention centers. The above speeds are those achieved in actual tests in 2020, and speeds are expected to increase during rollout.

1.7 CLOUD COMPUTING

This section provides a brief overview of cloud computing, which is dealt with in greater detail later in the book. Although the general concepts for cloud computing go back to the 1950s, cloud computing services first became available in the early 2000s, particularly targeted at large enterprises. Since then, cloud computing has spread to small- and medium-size businesses, and most recently to consumers. Apple's iCloud was launched in 2012 and had 20 million users within a week of launch. Evernote, the cloud-based note-taking and archiving service, launched in 2008, approached 100 million users in less than six years. In late 2014, Google announced that Google Drive had almost a quarter of a billion active users. Here, we look at the key elements of clouds, including cloud computing, cloud networking, and cloud storage.

1.7.1 Cloud Computing Concepts

There is an increasingly prominent trend in many organizations to move a substantial portion or even all IT operations to an Internet-connected infrastructure known as enterprise cloud computing. At the same time, individual users of PCs and mobile devices are relying more and more on cloud computing services to back up data, sync devices, and share, using personal cloud computing.

The National Institute of Standards and Technology (NIST) defines the essential characteristics of cloud computing as follows:

- Broad network access: Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (for example, mobile phones, laptops, and personal digital assistants [PDAs]) and other traditional or cloud-based software services.

- Rapid elasticity: Cloud computing enables you to expand and reduce resources according to your specific service requirement. For example, you may need many server resources for the duration of a specific task. You can then release these resources upon completion of the task.

- Measured service: Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (for example, storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.

- On-demand self-service: A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service provider. Because the service is on demand, the resources are not permanent parts of your IT infrastructure.

- Resource pooling: The provider’s computing resources are pooled to serve multiple consumers using a multitenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a degree of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources, but may be able to specify location at a higher level of abstraction (for example, country, state, or data center). Examples of resources include storage, processing, memory, network bandwidth, and virtual machines. Even private clouds tend to pool resources between different parts of the same organization.

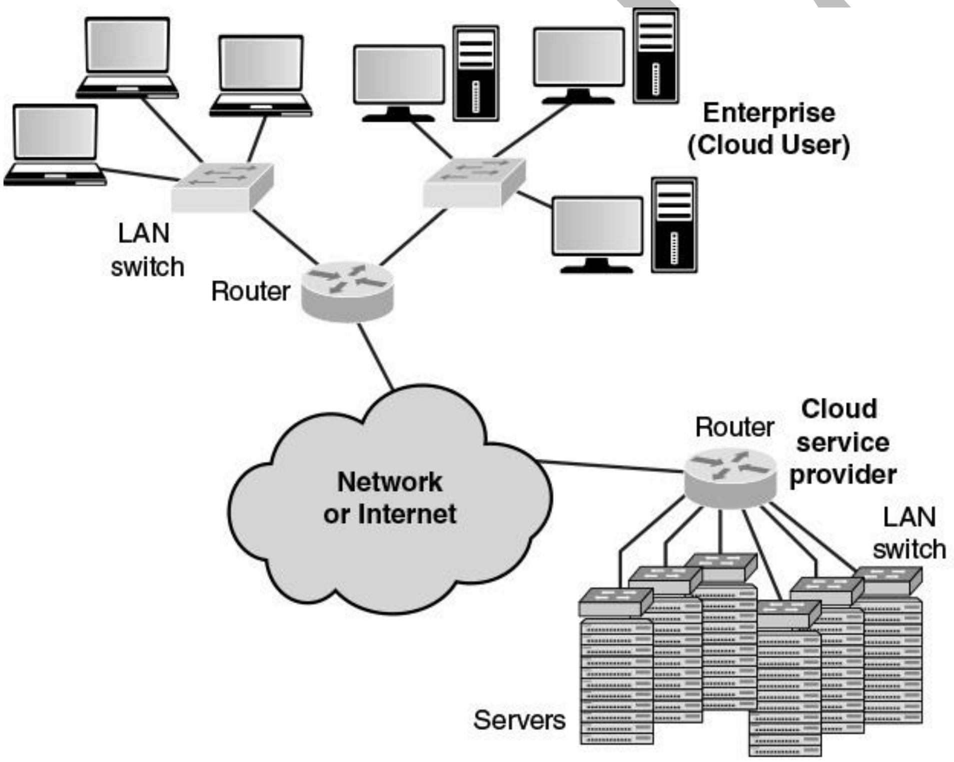

Figure 1.6 Cloud Computing Context

Figure 1.6 illustrates the typical cloud service context. An enterprise maintains workstations within an enterprise LAN or set of LANs, which are connected by a router through a network or the Internet to the cloud service provider. The cloud service provider maintains a massive collection of servers, which it manages with a variety of network management, redundancy, and security tools. In the figure, the cloud infrastructure is shown as a collection of blade servers, which is a common architecture.

1.7.2 The Benefits of Cloud Computing

Cloud computing benefits include –

a. Flexibility - Users can scale services to fit their needs, customize applications and access cloud services from anywhere with an internet connection. Flexibility further integrates:

   i. Scalability - Cloud infrastructure scales on demand to support fluctuating workloads.

   ii. Storage options - Users can choose public, private or hybrid storage offerings, depending on security needs and other considerations.

   iii. Control choices - Organizations can determine their level of control with as-a-service options. These include software as a service (SaaS), platform as a service (PaaS) and infrastructure as a service (IaaS).

   iv. Tool selection - Users can select from a menu of prebuilt tools and features to build a solution that fits their specific needs.

   v. Security features - Virtual private cloud, encryption and API keys help keep data secure.

b. Efficiency - Enterprise users can get applications to market quickly, without worrying about underlying infrastructure costs or maintenance. This could be further explained as:

   i. Accessibility - Cloud-based applications and data are accessible from virtually any internet-connected device.

   ii. Speed to market - Developing in the cloud enables users to get their applications to market quickly.

   iii. Data security - Hardware failures do not result in data loss because of networked backups.

   iv. Savings on equipment - Cloud computing uses remote resources, saving organizations the cost of servers and other equipment.

   v. Pay structure - A “utility” pay structure means users only pay for the resources they use.

c. Strategic value - Cloud services give enterprises a competitive advantage by providing the most innovative technology available consisting:

   i. Streamlined work - Cloud service providers (CSPs) manage underlying infrastructure, enabling organizations to focus on application development and other priorities.

   ii. Regular updates - Service providers regularly update offerings to give users the most up-to-date technology.

   iii. Collaboration - Worldwide access means teams can collaborate from widespread locations.

   iv. Competitive edge - Organizations can move more nimbly than competitors who must devote IT resources to managing infrastructure.

Cloud computing provides economies of scale, professional network management, and professional security management. These features can be attractive to companies large and small, government agencies, and individual PC and mobile users. The individual or company needs to pay only for the storage capacity and services they need.

The user, be it company or individual, does not have the hassle of setting up a database system, acquiring the hardware they need, doing maintenance, and backup up the data; all this is part of the cloud service. In theory, another big advantage of using cloud computing to store your data and share it with others is that the cloud provider takes care of security.

Alas, the customer is not always protected. There have been several security failures among cloud providers. Evernote made headlines in early 2013 when it told all its users to reset their passwords after an intrusion was discovered.

1.7.3 Cloud Networking

Cloud networking refers to the networks and network management functionality that must be in place to enable cloud computing. Many cloud computing solutions rely on the Internet, but that is only a piece of the networking infrastructure. One example of cloud networking is the provisioning high-performance/high-reliability networking between the provider and subscriber. In this case, some or all the traffic between an enterprise and the cloud bypasses the

Internet and uses dedicated private network facilities owned or leased by the cloud service provider.

More generally, cloud networking refers to the collection of network capabilities required to access a cloud, including making use of specialized services over the Internet, linking enterprise data centers to a cloud, and using firewalls and other network security devices at critical points to enforce access security policies.

1.7.4 Cloud Storage

We can think of cloud storage as a subset of cloud computing. In essence, cloud storage consists of database storage and database applications hosted remotely on cloud servers. Cloud storage enables small businesses and individual users to take advantage of data storage that scales with their needs and to take advantage of a variety of database applications without having to buy, maintain, and manage the storage assets.

1.8 INTERNET OF THINGS

The Internet of Things (IoT) is the latest development in the long and continuing revolution of computing and communications. Its size, ubiquity, and influence on everyday lives, business, and government dwarf any technical advance that has gone before. This section provides a brief overview of the IoT, which is dealt with in greater detail later in the book.

1.8.1 Things on the Internet of Things

The Internet of Things (IoT) is a term that refers to the expanding interconnection of smart devices, ranging from appliances to tiny sensors. A dominant theme is the embedding of short-range mobile transceivers into a wide array of gadgets and everyday items, enabling new forms of communication between people and things, and between things themselves.

The Internet now supports the interconnection of billions of industrial and personal objects, usually through cloud systems. The objects deliver sensor information, act on their environment, and in some cases modify themselves, to create overall management of a larger system, like a factory or city.

The IoT is primarily driven by deeply embedded devices. These devices are low-bandwidth, low-repetition data-capture and low-bandwidth data-usage appliances that communicate with each other and provide data via user interfaces.

Embedded appliances, such as high-resolution video security cameras, Video over IP (VoIP) phones, and a handful of others, require high bandwidth streaming capabilities. Yet countless products simply require packets of data to be intermittently delivered.

1.8.2 Evolution

With reference to end systems supported, the Internet has gone through roughly four generations of deployment culminating in IoT:

1. **Information technology (IT):** PCs, servers, routers, firewalls, and so on, bought as IT devices by enterprise IT people, primarily using wired connectivity.

2. **Operational technology (OT):** Machines/appliances with embedded IT built by non-IT companies, such as medical machinery, SCADA (supervisory control and data acquisition), process control, and kiosks, bought as appliances by enterprise OT people and primarily using wired connectivity.

3. **Personal technology:** Smartphones, tablets, and eBook readers bought as IT devices by consumers (employees) exclusively using wireless connectivity and often multiple forms of wireless connectivity.

4. **Sensor/actuator technology:** Single-purpose devices bought by consumers, IT, and OT people exclusively using wireless connectivity, generally of a single form, as part of larger systems.

It is the fourth generation that is usually thought of as the IoT, and which is marked using billions of embedded devices.

1.8.3 Layers of the Internet of Things

Both the business and technical literature often focus on two elements of the Internet of Things - the “things” that are connected, and the Internet that interconnects them. It is better to view the IoT as a massive system, which consists of five layers:

1. **Sensors and actuators**: Sensors observe their environment and report back quantitative measurements of such variables as temperature, humidity, presence, or absence of some observable, and so on. Actuators operate on their environment, such as changing a thermostat setting or operating a valve.

2. **Connectivity**: A device may connect via either a wireless or wired link into a network to send collected data to the

appropriate data center (sensor) or receive operational commands from a controller site (actuator).

3. **Capacity**: The network supporting the devices must be able to handle a potentially huge flow of data.

4. **Storage**: There needs to be a large storage facility to store and maintain backups of all the collected data. This is typically a cloud capability.

5. **Data analytics**: For large collections of devices, “big data” is generated, requiring a data analytics capability to process the data flow.

All these layers are essential to an effective use of the IoT concept.

1.9 NETWORK CONVERGENCE

Network convergence refers to the merger of previously distinct telephony and information technologies and markets. You can think of this convergence in terms of a three-layer model of enterprise communications:

- **Application convergence**: These are seen by the end users of a business. Convergence integrates communications applications, such as voice calling (telephone), voice mail, e-mail, and instant messaging, with business applications, such as workgroup collaboration, customer relationship management, and back-office functions. With convergence, applications provide rich features that incorporate voice, data, and video in a seamless, organized, and value-added manner. One example is multimedia messaging, which enables a user to use a single interface to access messages from a variety of sources (for example, office voice mail, e-mail, SMS text messages, and mobile voice mail).

- **Enterprise services:** At this level, the manager deals with the information network in terms of the services that must be available to ensure that users can take full advantage of the applications that they use. For example, network managers need to make sure that appropriate privacy mechanisms and authentication services are in place to support convergence-based applications. They may also be able to track user locations to support remote print services and network storage facilities for mobile workers. Enterprise network management services may also include setting up collaborative environments for various users, groups, and applications and QoS provision.

- **Infrastructure**: The network and communications infrastructure consist of the communication links, LANs, WANs, and Internet connections available to the enterprise. Increasingly, enterprise network infrastructure also includes private/public cloud connections to data centers that host high-volume data storage and web services. A key aspect of convergence at this level is the ability to carry voice, image, and video over networks that were originally designed to carry data traffic. Infrastructure convergence has also occurred for networks that were designed for voice traffic. For example, video, image, text, and data are routinely delivered to smartphone users over cell phone networks.

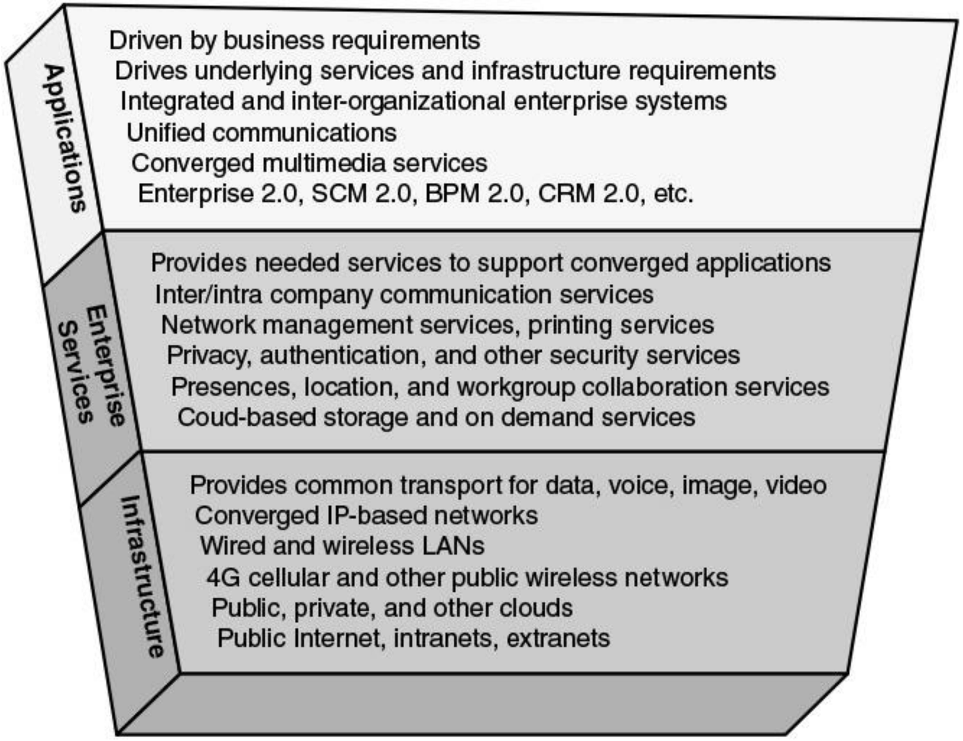
Applications
Driven by business requirements
Drives underlying services and infrastructure requirements
Integrated and inter-organizational enterprise systems
Unified communications
Converged multimedia services
Enterprise 2.0, SCM 2.0, BPM 2.0, CRM 2.0, etc.
Enterprise Services
Provides needed services to support converged applications
Inter/intra company communication services
Network management services, printing services
Privacy, authentication, and other security services
Presences, location, and workgroup collaboration services
Coud-based storage and on demand services
Infrastructure
Provides common transport for data, voice, image, video
Converged IP-based networks
Wired and wireless LANs
4G cellular and other public wireless networks
Public, private, and other clouds
Public Internet, intranets, extranets

Figure 1.7 Business-driven convergence

Figure 1.7 illustrates the major attributes of the three-layer model of enterprise communications. In simple terms, convergence involves moving an organization’s voice, video, and image traffic to a single network infrastructure. This often involves integrating distinct voice and data networks into a single network infrastructure and extending the infrastructure to support mobile users. The foundation of this convergence is packet-based transmission using the Internet Protocol (IP).

Using IP packets to deliver all varieties of communications traffic, sometimes referred to as everything over IP, enables the underlying infrastructure to deliver a wide range of useful applications to business users.

Convergence brings many benefits, including simplified network management, increased efficiency, and greater flexibility at the application level. For example, a converged network infrastructure provides a predictable platform on which to build new add applications that combine video, data, and voice. This makes it easier for developers to create innovative mashups and other value-added business applications and services.

The following list summarizes three key benefits of IP network convergence:

1. **Cost savings**: A converged network can provide significant double-digit percent reductions in network administration, maintenance, and operating costs; converging legacy networks onto a single IP network enables better use of existing resources, and implementation of centralized capacity planning, asset management, and policy management.

2. **Effectiveness**: The converged environment has the potential to provide users with great flexibility, irrespective of where they are. IP convergence allows companies to create a more mobile workforce. Mobile workers can use a virtual private network (VPN) to remotely access business applications and communication services on the corporate network. A VPN helps maintain enterprise network security by separating business traffic from other Internet traffic.

3. **Transformation**: Because they are modifiable and interoperable, converged IP networks can easily adapt to new functions and features as they become available through technological advancements without having to install new infrastructure. Convergence also enables the enterprise-wide adoption of global standards and best practices, thus providing better data, enhanced real-time decision making, and improved execution of key business processes and operations. The result is enhanced agility and innovation, the key ingredients of business innovation.

These compelling business benefits are motivating companies to invest in converged network infrastructures. Businesses, however, are keenly aware of the downside of convergence: having a single network means a single point of failure. Given their reliance on ICT (information and communications technology), today's converged enterprise network infrastructures typically include redundant components and back up systems to increase network resiliency and lessen the severity of network outages.

1.10 UNIFIED COMMUNICATIONS

While enterprise network convergence focuses on the consolidation of traditionally distinct voice, video, and data communications networks into a common infrastructure, Unified Communications (UC) focuses on the integration of real-time communication services to optimize business processes. As with converged enterprise networks, IP is the cornerstone on which UC systems are built.

Key elements of Unified Communications include the following:

1. UC systems typically provide a unified user interface and consistent user experience across multiple devices and media.
2. UC merges real-time communications services with non-real-time services and business process applications.

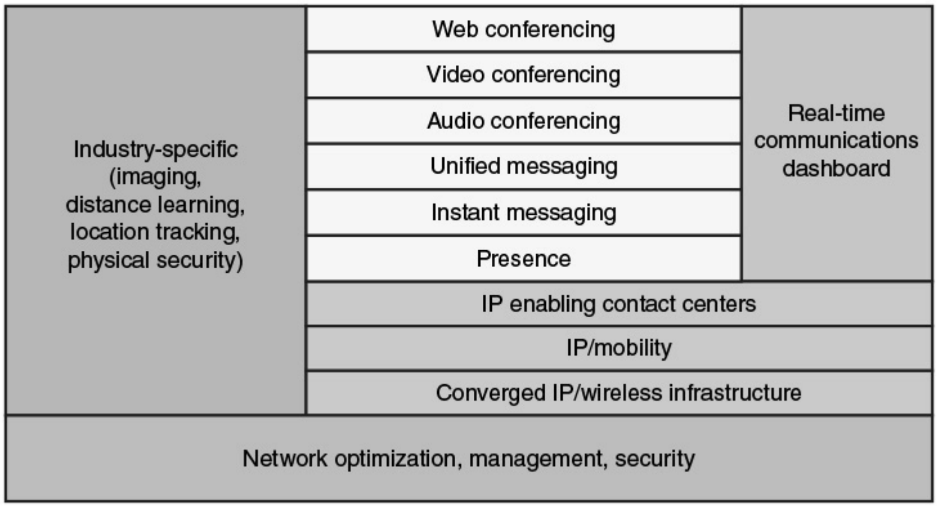
Industry-specific (imaging, distance learning, location tracking, physical security)
Web conferencing
Video conferencing
Audio conferencing
Unified messaging
Instant messaging
Presence
Real-time communications dashboard
IP enabling contact centers
IP/mobility
Converged IP/wireless infrastructure
Network optimization, management, security

Figure 1.8 Elements of a Unified Communications Architecture

Figure 1.8 shows the typical components of a UC architecture and how they relate to one another.

The key elements of this architecture are as follows:

- **Real-time communications (RTC) dashboard:** An RTC dashboard is a key component of UC architecture. This is the element that provides UC users with a unified user interface across communication devices. Ideally, the user has a consistent interface no matter what communication device the user is currently using, whether it is a cell phone, wireless tablet computer, desktop system, or office telephone attached to the corporate private branch exchange (PBX). As you can see in Figure 1.8, RTC dashboards provide access to real-time communication services such as instant messaging, audio and video conferencing, and interactive whiteboards; RTC dashboards also provide access to non-real-time services such as unified messaging (e-mail, voice mail, fax, and SMS) in unified view. An RTC dashboard includes presence information about co-workers and partners so that users can know on the fly which colleagues are available to communicate or join a collaborative communication session. RTC dashboards have become necessities in organizations that require high levels of communication and collaboration to support business processes.

- **Web conferencing:** Refers to live meetings or presentations in which participants access the meeting or presentation via a mobile device or the web, either over the Internet, or corporate intranet. Web conferences often include data sharing through web-connected interactive white boards (IWBs).

- **Audio conferencing:** Also called conference calling, refers to a live meeting in which participants are linked together for audio transmission and reception. A participant may be on a landline, mobile phone, or at a “softphone” - a computer equipped with microphone and speaker.

- **Unified messaging:** Unified messaging systems provide a common repository for messages from multiple sources. It allows users to retrieve saved e-mail, voice mail, and fax messages from a computer, telephone, or mobile device. Computer users can select and play voice-mail recordings that appear in their unified messaging inboxes. Telephone users can both retrieve voice mail and hear text-to-voice translations of e-mail messages. Messages of any type can be saved, answered, filed, sorted, and forwarded. Unified messaging systems relieve business users from having to monitor multiple voice mailboxes by enabling voicemail messages received by both office phones and cell phones to be saved to the same mailbox. With UC, users can use any device at any time to retrieve e-mail or voicemail from unified messaging mailboxes.

- **Instant messaging (IM):** Real-time text-based messaging between two or more participants. IM is like online chat because it is text-based and exchanged bidirectionally in real time. IM is distinct from chat in that IM clients use contact (or buddy) lists to facilitate connections between known users, whereas online chat can include text-based exchanges between anonymous users.

- **Video teleconferencing (VTC):** Videoconferencing allows users in two or more locations to interact simultaneously via two-way video and audio transmission. UC systems enable users to participate in video conferences via desktop computers, smartphones, and mobile devices.

- **Presence**: The capability to determine, in real time, where someone is, how that person prefers to be reached, and even what the person is currently doing. Presence information shows the individual's availability state before co-workers attempt to contact them person. It was once considered simply an underlying technology to instant messaging (for example, "available to chat" or "busy") but has been broadened to include whether co-workers are currently on office or mobile phones, logged in to a computer, involved in a video call or in a meeting, or out of the office for lunch or vacation. A co-worker's geographic location is becoming more common as an element in presence information for several business reasons, including the capability to quickly respond to customer emergencies. Business has embraced presence information because it facilitates more efficient and effective communication. It helps eliminate inefficiencies associated with "phone tag" or composing and sending e-mails to someone who could more quickly answer a question over the phone or with a quick meeting.

- **IP enabling contact centers:** Refers to the use of IP-based unified communications to enhance customer contact center functionality and performance. The unified communications infrastructure makes use of presence technology to enable customers and internal enterprise employees to be quickly connected to the required expert or support person.

In addition, this technology supports mobility, so that call center personnel need not be located at a particular office or remain in a particular place. Finally, the UC infrastructure enables the call center employee to quickly access other employees and information assets, including data, video, image, and audio.

- **IP/mobility:** Refers to the delivery of information to and collection of information from enterprise personnel who are usually mobile, using an IP network infrastructure. In a typical enterprise, upward of 30 percent of employees use some form of weekly remote access technology in the performance of their jobs.

- **Converged IP/wireless infrastructure:** A unified networking and communications-based IP packet transfer to support voice, data, and video transmission and can be extended to include local- and wide-area wireless communications. UC-enabled mobile devices can switch between Wi-Fi and cellular systems in the middle of a communication session.

  For example, a UC user could receive a co-worker's call via a smartphone connected to Wi-Fi network at home, continue the conversation while driving to work over a cellular network connection, and could end the call at the office while connected to the business's Wi-Fi network. Both handoffs (home Wi-Fi to cellular and cellular to office Wi-Fi) would take place seamlessly and transparently without dropping the call.

The importance of UC is not only that it integrates communication channels but also that it offers a way to integrate communication functions and business applications. Three major categories of benefits are typically realized by organizations that use UC:

- **Personal productivity gains:** Presence information helps employees find each other and choose the most effective way to communicate in real time. Less time is wasted calling multiple numbers to locate co-workers or checking multiple worked-related voice mailboxes. Calls from VIP contacts can be routed simultaneously to all a UC user's phone devices (office phone, softphone, smartphone, home phone) to ensure faster responsiveness to customers, partners, and co-workers. With mobile presence information capabilities, employees who are geographically closest can be dispatched to address a problem.

- **Workgroup performance gains:** UC systems support real-time collaboration among team members, which facilitates workgroup performance improvements. Examples include the use of presence information to speed identification of an available individual with the right skills a work team needs to address a problem. Enhanced conferencing capabilities with desktop VTC and interactive white boards and automated business rules to route or escalate communications also help to increase workgroup performance.

- **Enterprise-level process improvements:** IP convergence enables UC to be integrated with enterprise-wide and departmental-level applications, business processes, and workflows. UC-enabled enhanced communications with customers, suppliers, and business partners are redefining best practices for customer relationship management (CRM), supply chain management (SCM), and other enterprise-wide applications and are transforming relationships among members of business networks. Communication-enabled business processes (CEBP) are fueling competition in several industries, including financial services, healthcare, and retail.

1.11 SUMMARY

- Users connect to network-based services and content through a wide variety of network access facilities. These include digital subscriber line (DSL) and cable modems, Wi-Fi, and Worldwide Interoperability for Microwave Access (WiMAX) wireless modems, and cellular modems.

- At the periphery of an IP backbone are routers that provide connectivity to external networks and users.

- Some of the core routers may be purely internal, providing redundancy and additional capacity without serving as edge routers.

- Two recent extensions of Ethernet technology have enhanced and broadened the use of Ethernet in the home: powerline carrier (PLC) and Power over Ethernet (PoE).

- A tremendous advantage of Ethernet is that it is possible to scale the network, both in terms of distance and data rate, with the same Ethernet protocol and associated quality of service (QoS) and security standards.

- The technology also allows the construction of MANs and WANs that connect geographically dispersed LANs between campuses or points of presence (PoPs).

- In the consumer and residential market, IEEE 802.11ad is likely to be popular as a low-power, short-distance wireless LAN capability with little likelihood of interfering with other devices.

- Cloud networking refers to the networks and network management functionality that must be in place to enable cloud computing.

- Convergence brings many benefits, including simplified network management, increased efficiency, and greater flexibility at the application level.

- UC systems typically provide a unified user interface and consistent user experience across multiple devices and media.

- The importance of UC is not only that it integrates communication channels but also that it offers a way to integrate communication functions and business applications.

- Cloud computing benefits include flexibility, efficiency and strategic value.

1.12 REVIEW QUESTION

1. Explain the key elements and their relationships of a modern networking ecosystem, including end users, network providers, application providers and application service providers.

2. Discuss the motivation for the typical network hierarchy of access networks, distribution networks, and core networks.

3. Present an overview of Ethernet, including a discussion of its application areas and common data rates.

4. Present an overview of Wi-Fi, including a discussion of its application areas and common data rates.

5. Understand the differences between the five generations of cellular networks.

6. Present an overview of cloud computing concepts.

7. Describe the Internet of Things.

8. Explain the concepts of network convergence and unified communications.

1.13 REFERENCES

1. Foundations of Modern Networking: SDN, NFV, QoE, IoT, and Cloud, First printing: November 2015 by William Stallings, Copyright © 2016 by Pearson Education, Inc. ISBN-13: 978-0-13-417539-3

2. Research paper on “Use of Ethernet Technology in Computer Network” by Zobair Ullah, published in Global Journal of Computer Science and Technology Network, Web & Security Volume 12 Issue 14 Version 1.0 Year 2012, available online at https://core.ac.uk/download/pdf/231149472.pdf

3. IEEE 802.3 standards working group’s information available online at https://en.wikipedia.org/wiki/IEEE_802.3

4. Note on 5G available online at https://en.wikipedia.org/wiki/5G

5. Online information available regarding benefits of cloud computing - https://www.ibm.com/in-en/cloud/learn/benefits-of-cloud-computing

6. Online article “The future of networking: A guide to the intelligent network” by Lauren Horwitz, available on Cisco Digital Network Architecture (DNA) portal online at https://www.cisco.com/c/en/us/solutions/enterprise-networks/future-of-networking.html

❖❖❖❖

2

REQUIREMENTS AND TECHNOLOGY

Unit Structure:

2.1 OBJECTIVES

After studying this chapter, you should be able to:

- Present an overview of the major categories of packet traffic on the Internet and internets, including elastic, inelastic, and real-time traffic.
- Discuss the traffic demands placed on contemporary networks by big data, cloud computing, and mobile traffic.

- Explain the concept of quality of service.
- Explain the concept of quality of experience.
- Understand the essential elements of routing.
- Understand the effects of congestion and the types of techniques used for congestion control.
- Compare and contrast software-defined networking and network functions virtualization.

2.2 TYPES OF NETWORK AND INTERNET TRAFFIC

Most of the Internet traffic today is generated by traditional data applications; such traffic is for the most part burst and is well served by the best-effort service that the Internet provides. With the growth and ubiquity of the Internet witnessed in recent years, new applications are being contemplated, introducing new traffic types and new requirements, which in turn require new services from the network which cater to these characteristics and requirements. Furthermore, as the Internet becomes a network on which many businesses rely, it becomes crucial for the network response time to be unaffected by increases in the load on the network.

The phrase **traffic classification** is used to describe methods of classifying traffic based on features passively observed in the traffic, and according to specific classification goals. One might only have a coarse classification goal, i.e., whether its transaction-oriented, bulk-transfer, or peer-to-peer file sharing. Or one might have a finer-grained classification goal, i.e., the exact application represented by the traffic. Traffic features could include the port number, application payload, or temporal, packet size, and addressing characteristics of the traffic. Methods to classify include exact matching, e.g., of port number or payload, heuristic, or machine learning (statistics).

Traffic on the Internet and enterprise networks can be divided into two broad categories: elastic and inelastic. A consideration of their differing requirements clarifies the need for an enhanced networking architecture.

Elastic Traffic

Elastic traffic is that which can adjust, over wide ranges, to changes in delay and throughput across an internet and still meet the needs of its applications. This is the traditional type of traffic supported on TCP/IP-based internets and is the type of traffic for which internets were designed. Applications that generate such traffic typically use Transmission Control Protocol (TCP) or User Datagram Protocol

(UDP) as a transport protocol. In the case of UDP, the application will use as much capacity as is available up to the rate that the application generates data. In the case of TCP, the application will use as much capacity as is available up to the maximum rate that the end-to-end receiver can accept data. Also, with TCP, traffic on individual connections adjusts to congestion by reducing the rate at which data are presented to the network. Applications that can be classified as elastic include the common applications that operate over TCP or UDP, including file transfer (File Transfer Protocol / Secure FTP [FTP/SFTP]), electronic mail (Simple Mail Transport Protocol [SMTP]), remote login (Telnet, Secure Shell [SSH]), network management (Simple Network Management Protocol [SNMP]), and web access (Hypertext Transfer Protocol / HTTP Secure [HTTP/HTTPS]). However, there are differences among the requirements of these applications, including the following:

- E-mail is generally insensitive to changes in delay.
- When file transfer is done via user command rather than as an automated background task, the user expects the delay to be proportional to the file size and so is sensitive to changes in throughput.
- With network management, delay is generally not a serious concern. However, if failures in an internet are the cause of congestion, then the need for SNMP messages to get through with minimum delay increases with increased congestion.
- Interactive applications, such as remote logon and web access, are sensitive to delay.

It is important to realize that it is not per-packet delay that is the quantity of interest. Observation of real delays across the Internet suggest that wide variations in delay do not occur. Because of the congestion control mechanisms in TCP, when congestion develops, delays only increase modestly before the arrival rate from the various TCP connections slow down. Instead, the quality of service (QoS) perceived by the user relates to the total elapsed time to transfer an element of the current application. For an interactive Telnet-based application, the element may be a single keystroke or single line. For web access, the element is a web page, which could be as little as a few kilobytes or could be substantially larger for an image-rich page.

For a scientific application, the element could be many megabytes of data. For very small elements, the total elapsed time is dominated by the delay time across the Internet. However, for larger elements, the total elapsed time is dictated by the sliding-window performance of TCP and is therefore dominated by the throughput achieved over the TCP connection. Thus, for large transfers, the transfer time is proportional to the size of the file and the degree to which the source slows because of congestion.

It should be clear that even if you confine your attention to elastic traffic, some service prioritizing and controlling traffic could be of benefit. Without such a service, routers are dealing evenhandedly with arriving IP packets, with no concern for the type of application and whether a particular packet is part of a large transfer element or a small one. Under such circumstances, and if congestion develops, it is unlikely that resources will be allocated in such a way as to meet the needs of all applications fairly. When inelastic traffic is added to the mix, the results are even more unsatisfactory.

Inelastic Traffic

Inelastic traffic does not easily adapt, if at all, to changes in delay and throughput across an internet. Examples of inelastic traffic include multimedia transmission, such as voice and video, and high-volume interactive traffic, such as an interactive simulation application (for example, airline pilot simulation). The requirements for inelastic traffic may include the following:

- Throughput: A minimum throughput value may be required. Unlike most elastic traffic, which can continue to deliver data with perhaps degraded service, many inelastic applications absolutely require a given minimum throughput.

- Delay: Also called latency. An example of a delay-sensitive application is stock trading; someone who consistently receives later service will consistently act later, and with greater disadvantage.

- Delay jitter: The magnitude of delay variation, called delay jitter, or simply jitter, is a critical factor in real-time applications. Because of the variable delay imposed by an internet, the interarrival times between packets are not maintained at a fixed interval at the destination. To compensate for this, the incoming packets are buffered, delayed sufficiently to compensate for the jitter, and then released at a constant rate to the software that is expecting a steady real-time stream. The larger the allowable delay variation, the longer the real delay in delivering the data and the greater the size of the delay buffer required at receivers. Real-time interactive applications, such as teleconferencing, may require a reasonable upper bound on jitter.

- Packet loss: Real-time applications vary in the amount of packet loss, if any, that they can sustain.

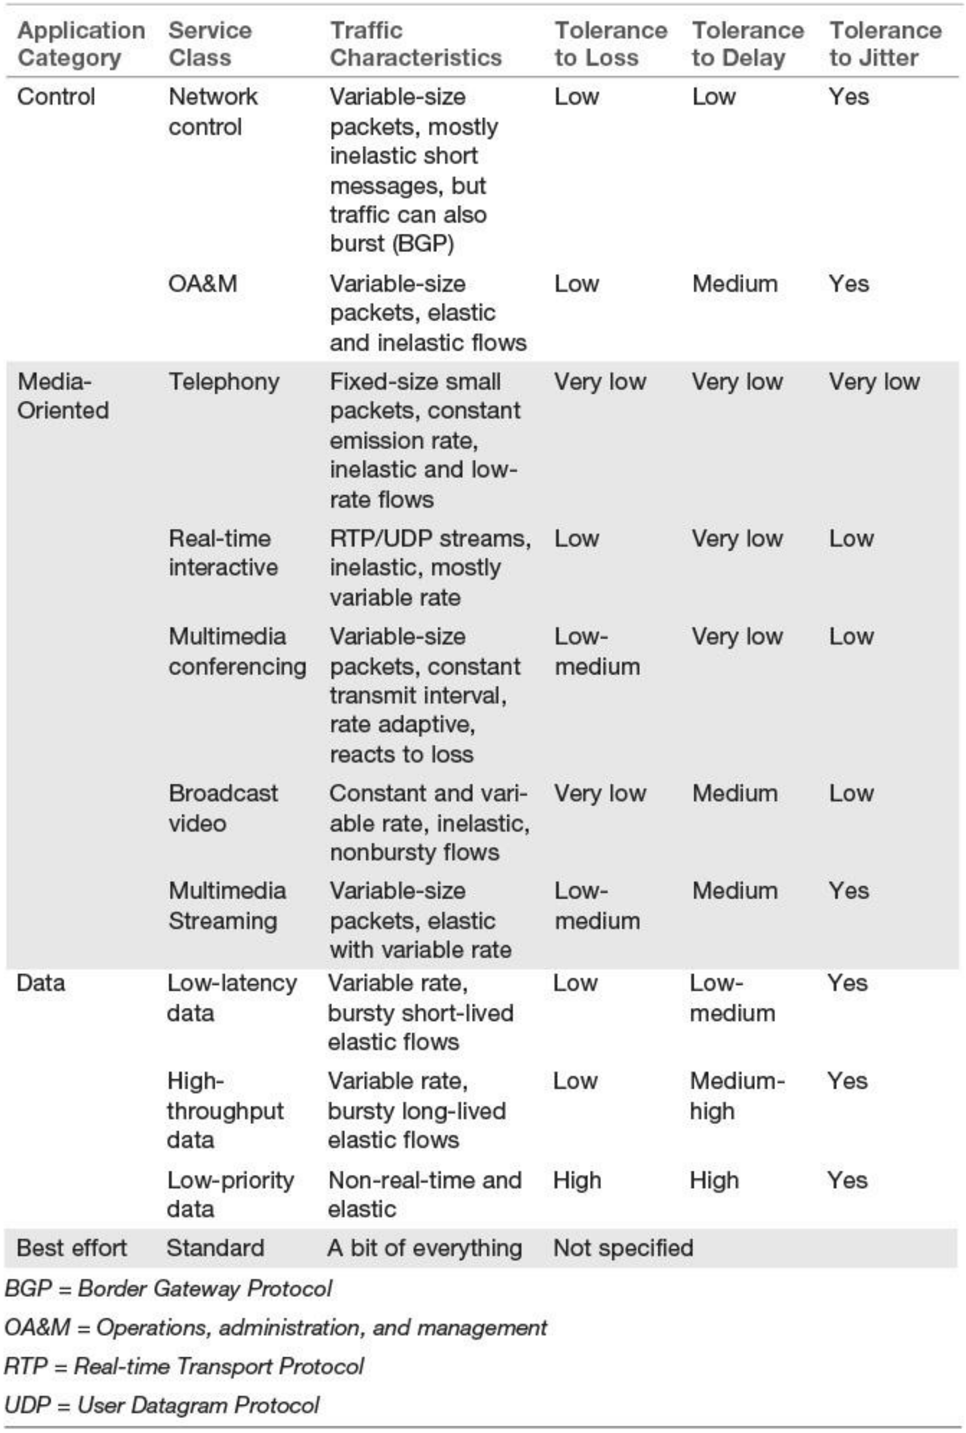

Application Category	Service Class	Traffic Characteristics	Tolerance to Loss	Tolerance to Delay	Tolerance to Jitter
Control	Network control	Variable-size packets, mostly inelastic short messages, but traffic can also burst (BGP)	Low	Low	Yes
	OA&M	Variable-size packets, elastic and inelastic flows	Low	Medium	Yes
Media-Oriented	Telephony	Fixed-size small packets, constant emission rate, inelastic and low-rate flows	Very low	Very low	Very low
	Real-time interactive	RTP/UDP streams, inelastic, mostly variable rate	Low	Very low	Low
	Multimedia conferencing	Variable-size packets, constant transmit interval, rate adaptive, reacts to loss	Low-medium	Very low	Low
	Broadcast video	Constant and variable rate, inelastic, nonbursty flows	Very low	Medium	Low
	Multimedia Streaming	Variable-size packets, elastic with variable rate	Low-medium	Medium	Yes
Data	Low-latency data	Variable rate, bursty short-lived elastic flows	Low	Low-medium	Yes
	High-throughput data	Variable rate, bursty long-lived elastic flows	Low	Medium-high	Yes
	Low-priority data	Non-real-time and elastic	High	High	Yes
Best effort	Standard	A bit of everything	Not specified		

BGP = Border Gateway Protocol

OA&M = Operations, administration, and management

RTP = Real-time Transport Protocol

UDP = User Datagram Protocol

Table 2.1 Service Class Characteristics

Table 2.1 above shows the loss, delay, and jitter characteristics of various classes of traffic, as specified in RFC 4594 (Configuration Guidelines for DiffServ Service Classes, August 2006).

Table 2.2 below gives examples of QoS requirements for various media-oriented applications [SZIG14]

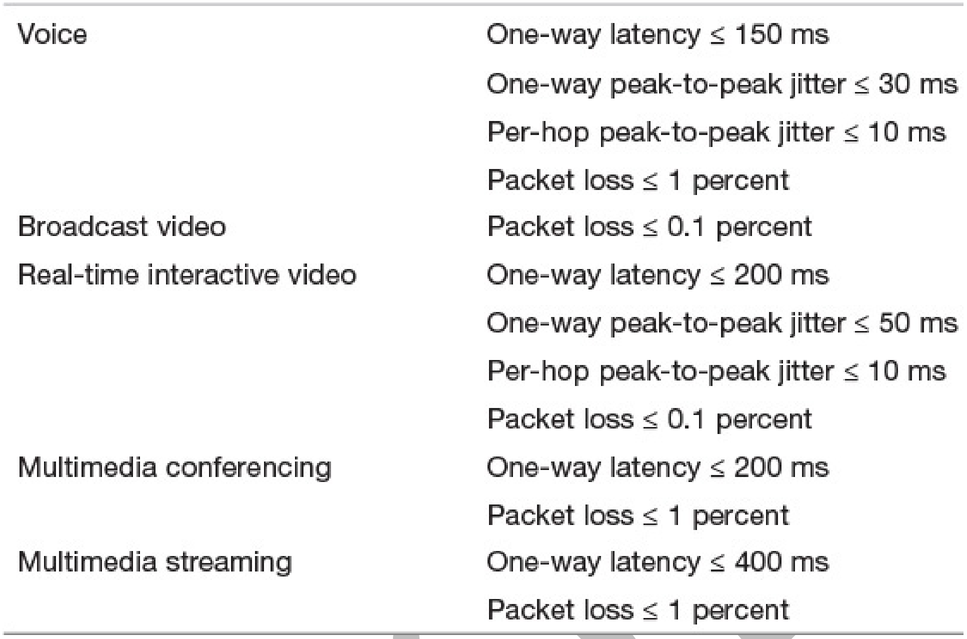

Voice	One-way latency ≤ 150 ms
	One-way peak-to-peak jitter ≤ 30 ms
	Per-hop peak-to-peak jitter ≤ 10 ms
	Packet loss ≤ 1 percent
Broadcast video	Packet loss ≤ 0.1 percent
Real-time interactive video	One-way latency ≤ 200 ms
	One-way peak-to-peak jitter ≤ 50 ms
	Per-hop peak-to-peak jitter ≤ 10 ms
	Packet loss ≤ 0.1 percent
Multimedia conferencing	One-way latency ≤ 200 ms
	Packet loss ≤ 1 percent
Multimedia streaming	One-way latency ≤ 400 ms
	Packet loss ≤ 1 percent

Table 2.2 QoS Requirements by Application Class

These requirements are difficult to meet in an environment with variable queuing delays and congestion losses. Accordingly, inelastic traffic introduces two new requirements into the internet architecture. First, some means is needed to give preferential treatment to applications with more demanding requirements. Applications need to be able to state their requirements, either ahead of time in some sort of service request function, or on the fly, by means of fields in the IP packet header. The former approach provides more flexibility in stating requirements, and it enables the network to anticipate demands and deny new requests if the required resources are unavailable. This approach implies the use of some sort of resource reservation protocol.

An additional requirement in supporting inelastic traffic in an internet architecture is that elastic traffic must still be supported. Inelastic applications typically do not back off and reduce demand in the face of congestion, in contrast to TCP-based applications. Therefore, in times of congestion, inelastic traffic will continue to supply a high load, and elastic traffic will be crowded off the internet. A reservation protocol can help control this situation by denying service requests that would leave too few resources available to handle current elastic traffic.

2.2.1 Reat-Time Traffic Characteristics

As mentioned, a common example of inelastic traffic is real-time traffic. With traditional elastic applications, such as file transfer, electronic mail, and client/server applications including the web, the performance metrics of interest are generally throughput and delay. There is also a concern with reliability, and mechanisms are used to make sure that no data are lost, corrupted, or misordered during transit. By contrast, real-time applications are concerned with timing issues as well as packet loss. In most cases, there is a requirement that data be delivered at a constant rate equal to the sending rate. In other cases, a deadline is associated with each block of data, such that the data are not usable after the deadline has expired.

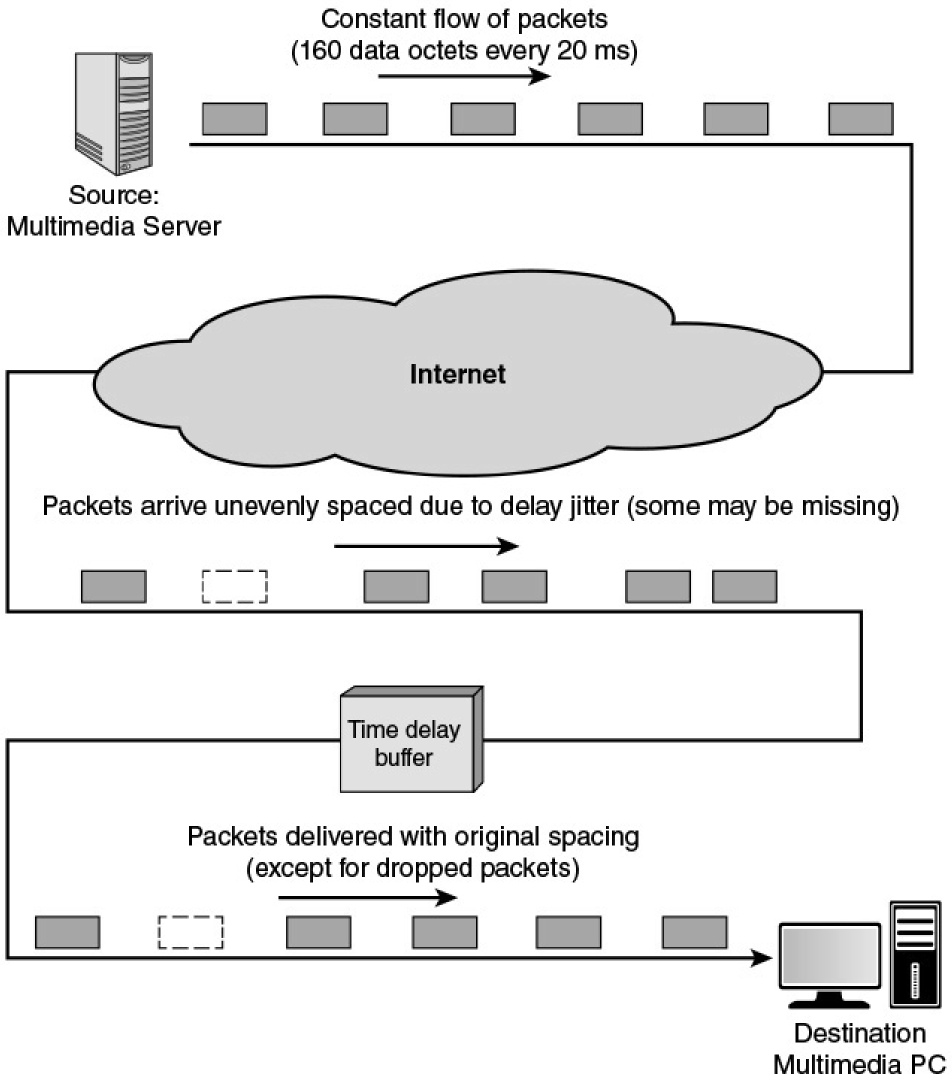
Constant flow of packets
(160 data octets every 20 ms)
Source:
Multimedia Server
Internet
Packets arrive unevenly spaced due to delay jitter (some may be missing)
Time delay buffer
Packets delivered with original spacing
(except for dropped packets)
Destination
Multimedia PC

Figure 2.1 Real-Time Traffic

Figure 2.1 illustrates a typical real-time environment. Here, a server is generating audio to be transmitted at 64 kbps. The digitized audio is transmitted in packets containing 160 octets of data, so that one packet is issued every 20 ms. These packets are passed through an internet and delivered to a multimedia PC, which plays the audio in real time as it arrives. However, because of the variable delay imposed by the internet, the interarrival times between packets are not maintained at a fixed 20 ms at the destination. To compensate for this, the incoming packets are buffered, delayed slightly, and then released at a constant rate to the software that generates the audio. The buffer may be internal to the destination machine or in an external network device.

The compensation provided by the delay buffer is limited. For example, if the minimum end-to-end delay seen by any packet is 1 ms and the maximum is 6 ms, the delay jitter is 5 ms. As long as the time delay buffer delays incoming packets by at least 5 ms, the output of the buffer will include all incoming packets. However, if the buffer delayed packets by only 4 ms, any incoming packets that had experienced a relative delay of more than 4 ms (an absolute delay of more than 5 ms) would have to be discarded so as not to be played back out of order.

The description of real-time traffic so far implies a series of equal-size packets generated at a constant rate. This is not always the profile of the traffic. Figure 2.2 illustrates some of the common possibilities, as described in the list that follows.

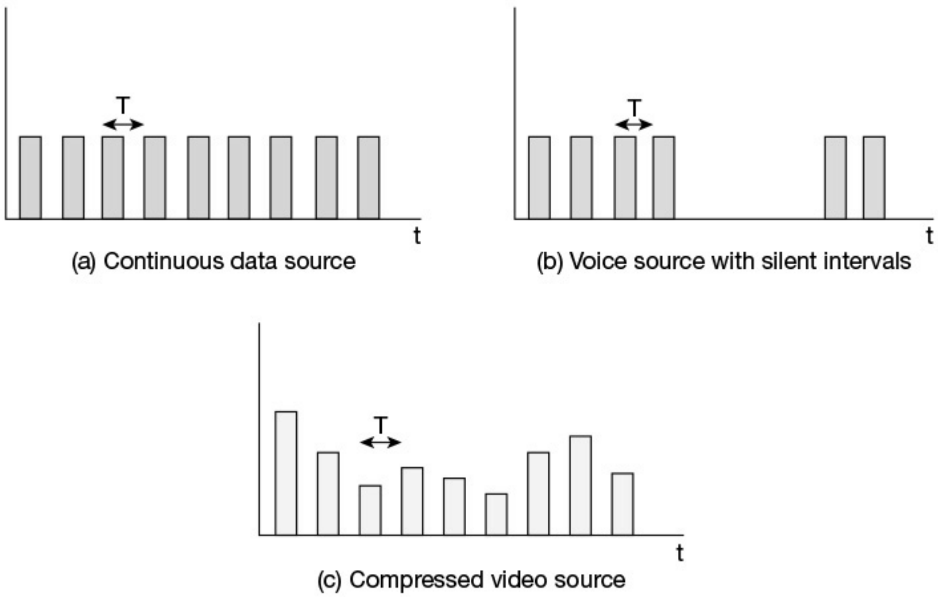

(a) Continuous data source

(b) Voice source with silent intervals

(c) Compressed video source

Figure 2.2 Real-Time Packet Transmission

- Continuous data source: Fixed-size packets are generated at fixed intervals. This characterizes applications that are constantly generating data, have few redundancies, and that are too important to compress in a lossy way. Examples are air traffic control radar and real-time simulations.

- On/off source: The source alternates between periods when fixed-size packets are generated at fixed intervals and periods of inactivity. A voice source, such as in telephony or audio conferencing, fits this profile.

- Variable packet size: The source generates variable-length packets at uniform intervals. An example is digitized video in which different frames may experience different compression ratios for the same output quality level.

2.3 DEMAND: BIG DATA, CLOUD COMPUTING, AND MOBILE TRAFFIC

Having looked at the types of traffic presented to the Internet and other IP-based networks, consider the application areas that are generating the greatest stress on network resources and management. Three areas stand out: big data, cloud computing, and mobility. All of these areas suggest the need for using powerful tools such as software-defined networking (SDN) and network functions virtualization (NFV) for network operation and management, and for using comprehensive QoS and quality of experience (QoE) systems for effective delivery of services over IP-based networks.

2.3.1 Big Data

In simple terms, big data refers to everything that enables an organization to create, manipulate, and manage very large data sets (measured in terabytes, petabytes, exabytes, and so on) and the facilities in which these are stored. Distributed data centers, data warehouses, and cloud-based storage are common aspects of today’s enterprise networks. Many factors have contributed to the merging of “big data” and business networks, including continuing declines in storage costs, the maturation of data mining and business intelligence (BI) tools, and government regulations and court cases that have caused organizations to stockpile large masses of structured and unstructured data, including documents, e-mail messages, voice-mail messages, text messages, and social media data. Other data sources being captured, transmitted, and stored include web logs, Internet documents, Internet search indexing, call detail records, scientific research data and results, military surveillance, medical records, video archives, and e-commerce transactions.

Data sets continue to grow with more and more being gathered by remote sensors, mobile devices, cameras, microphones, radio frequency identification (RFID) readers, and similar technologies. One study from a few years ago estimated that 2.5 exabytes (2.5 × 10 bytes) of data are created each day, and 90 percent of the data

in the world was created in the past two years. Those numbers are likely higher today.

Big Data Infrastructure Considerations

Traditional business data storage and management technologies include relational database management systems (RDBMS), network-attached storage (NAS), storage-area networks (SANs), data warehouses (DWs), and business intelligence (BI) analytics. Traditional data warehouse and BI analytics systems tend to be highly centralized within an enterprise infrastructure. These often include a central data repository with a RDBMS, high-performance storage, and analytics software, such as online analytical processing (OLAP) tools for mining and visualizing data.

Increasingly, big data applications are becoming a source of competitive value for businesses, especially those that aspire to build data products and services to profit from the huge volumes of data that they capture and store. There is every indication that the exploitation of data will become increasingly important to enterprises in the years ahead as more and more businesses reap the benefits of big data applications.

Big Data Networking Example

Key elements within the enterprise include the following:

- Data warehouse: The DW holds integrated data from multiple data sources, used for reporting and data analysis.

- Data management servers: Large banks of servers serve multiple functions with respect to big data. The servers run data analysis applications, such as data integration tools and analytics tools. Other applications integrate and structure data from enterprise activity, such as financial data, point-of-sale data, and e-commerce activity.

- Workstations / data processing systems: Other systems involved in the use of big data applications and in the generation of input to big data warehouses.

- Network management server: One or more servers responsible for network management, control, and monitoring.

Not shown in Figure 2.3 are other important network devices, including firewalls, intrusion detection/prevention systems (IDS/IPS), LAN switches, and routers.

The enterprise network can involve multiple sites distributed regionally, nationally, or globally. In addition, depending on the nature of the big data system, an enterprise can receive data from other enterprise servers, from dispersed sensors and other devices in an Internet of Things, in addition to multimedia content from content delivery networks.

The networking environment for big data is complex. The impact of big data on an enterprise's networking infrastructure is driven by the so-called three V's:

- Volume (growing amounts of data)
- Velocity (increasing speed in storing and reading data)
- Variability (growing number of data types and sources)

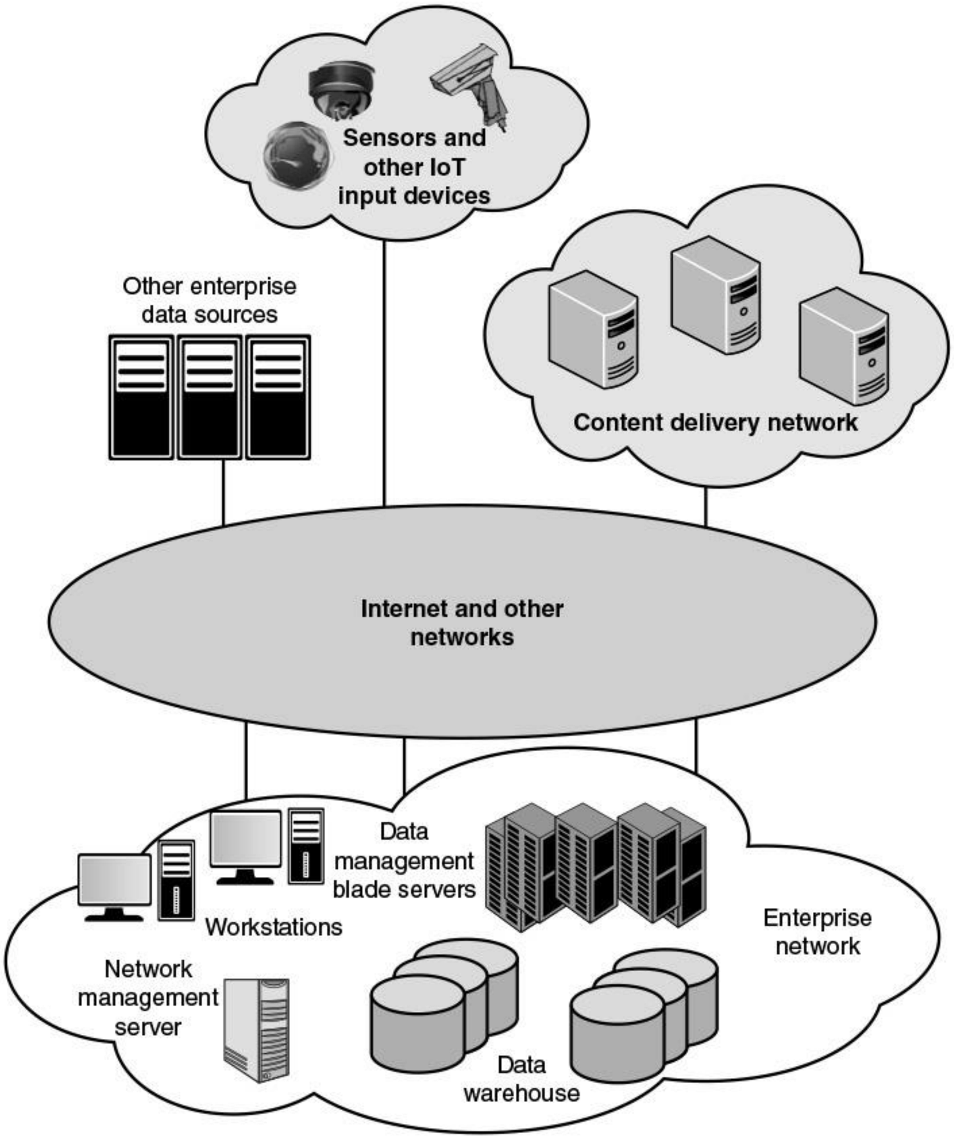
Sensors and other IoT input devices
Other enterprise data sources
Content delivery network
Internet and other networks
Data management blade servers
Workstations
Enterprise network
Network management server
Data warehouse

Figure 2.3 Big Data Networking Ecosystem

2.3.2 Cloud Computing

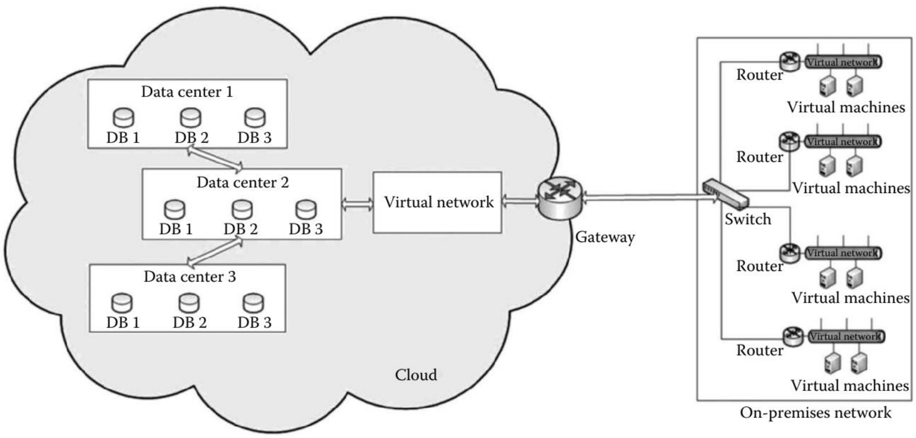

Figure 2.4 Cloud-based Network

A cloud-based network is an enterprise network that can be extended to the cloud shown in Figure 2.4. The cloud-based network allows an enterprise to distribute its network around the world. The cloud significantly simplifies the development of an enterprise network system. In the cloud, the underlying network is constructed by a cloud provider. All an enterprise needs to do is to connect its on-premises network to the network built in the cloud to form a global enterprise-class network system.

There is no initial capital investment in this type of global network system. Unlike the Internet, the cloud-based network provides centralized control over network visibility. Through the cloud-based network, the enterprise can provide a multitenant application, which is a software application that serves multiple tenants. Each tenant subscribes an instance of the application. Each tenant's data are isolated and remain invisible to other tenants. On the other hand, the maintenance and update of the application can be greatly simplified. The cloud-based network enables the enterprise to deploy IT infrastructures to remote locations in minutes.

The cloud-based network targets organizations with many sites around the world. There could be a couple of hundred to ten thousand employees working in multiple sites such as branch offices, schools in a school district, clinics, manufacturing facilities, or retail stores. Through the management tools deployed in the cloud, network administrators can manage the enterprise-distributed networks anywhere and anytime. The management tools can be used to manage cloud-hosted virtual machines and mobile services. They are used to accomplish tasks such as centralized management, remote monitoring, remote software and app installation, remote wiping, and security auditing.

There are three different major implementations of cloud computing. How organizations are using cloud computing is quite different at a granular level, but the uses generally fall into one of these three solutions.

Compute Clouds

Compute clouds allow access to highly scalable, inexpensive, on-demand computing resources that run the code that they are given. Three examples of compute clouds are

- Amazon's EC2
- Google App Engine

- Berkeley Open Infrastructure for Network Computing (BOINC)

Compute clouds are the most flexible in their offerings and can be used for sundry purposes; it simply depends on the application the user wants to access. You could close this book right now, sign up for a cloud computing account, and get started right away. These applications are good for any size organization, but large organizations might be at a disadvantage because these applications do not offer the standard management, monitoring, and governance capabilities that these organizations are used to. Enterprises are not shut out, however. Amazon offers enterprise-class support and there are emerging sets of cloud offerings like Terremark's Enterprise Cloud, which are meant for enterprise use.

Cloud Storage

One of the first cloud offerings was cloud storage and it remains a popular solution. Cloud storage is a big world. There are already more than 100 vendors offering cloud storage. This is an ideal solution if you want to maintain files off-site. Security and cost are the top issues in this field and vary greatly, depending on the vendor you choose. Currently, Amazon's S3 is the top player.

Cloud Applications

Cloud applications differ from compute clouds in that they utilize software applications that rely on cloud infrastructure. Cloud applications are versions of Software as a Service (SaaS) and include such things as web applications that are delivered to users via a browser or application like Microsoft Online Services. These applications offload hosting and IT management to the cloud.

Cloud applications often eliminate the need to install and run the application on the customer's own computer, thus alleviating the burden of software maintenance, ongoing operation, and support. Some cloud applications include

- Peer-to-peer computing (like Skype)
- Web applications (like MySpace or YouTube)
- SaaS (like Google Apps)
- Software plus services (like Microsoft Online Services)

2.3.3 Mobile Traffic

Following the extraordinary peak in traffic growth seen in 2018 and the first part of 2019, the growth rate has returned to a more normal level. The quarter-on-quarter growth for Q1 2020 was 14 percent. A change in consumer behavior caused by COVID-19 lockdown restrictions impacted mobile networks by geographically shifting traffic loads; for example, daytime loads moved, to a degree, from city centers to suburban residential areas due to home-working guidance. This effect was most pronounced in areas with limited penetration of fixed residential broadband connections. Generally, the traffic volumes were only modestly affected in mobile networks in markets where fixed network connections are common. Over the long-term, traffic growth is driven by both the rising number of smartphone subscriptions and an increasing average data volume per subscription, fueled primarily by more viewing of video content.

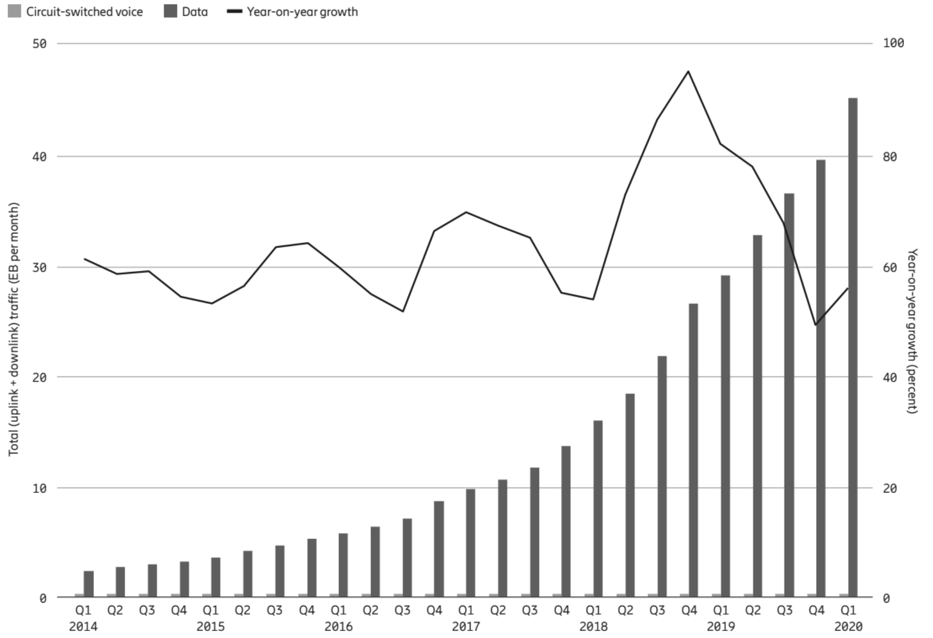

Figure 2.5: Global mobile network data traffic and year-on-year growth (exabytes per month)

Figure 2.5 shows total global monthly network data and voice traffic from Q1 2014 to Q1 2020, along with the year-on-year percentage change for mobile network data traffic. Mobile network data traffic depicted in Figure 2.5 also includes traffic generated by fixed wireless access (FWA) services and does not include DVB-H, Wi-Fi, or Mobile WiMAX. VoIP is included.

Global total mobile data traffic reached around 33 EB per month by the end of 2019 and is projected to grow by a factor close to 5 to reach 164 EB per month in 2025. Figure 2.6 represents the mobile data that will be consumed by over 6 billion people using smartphones, laptops, and a multitude of new devices at that time. (This graph does not include traffic generated by fixed wireless access (FWA) services)

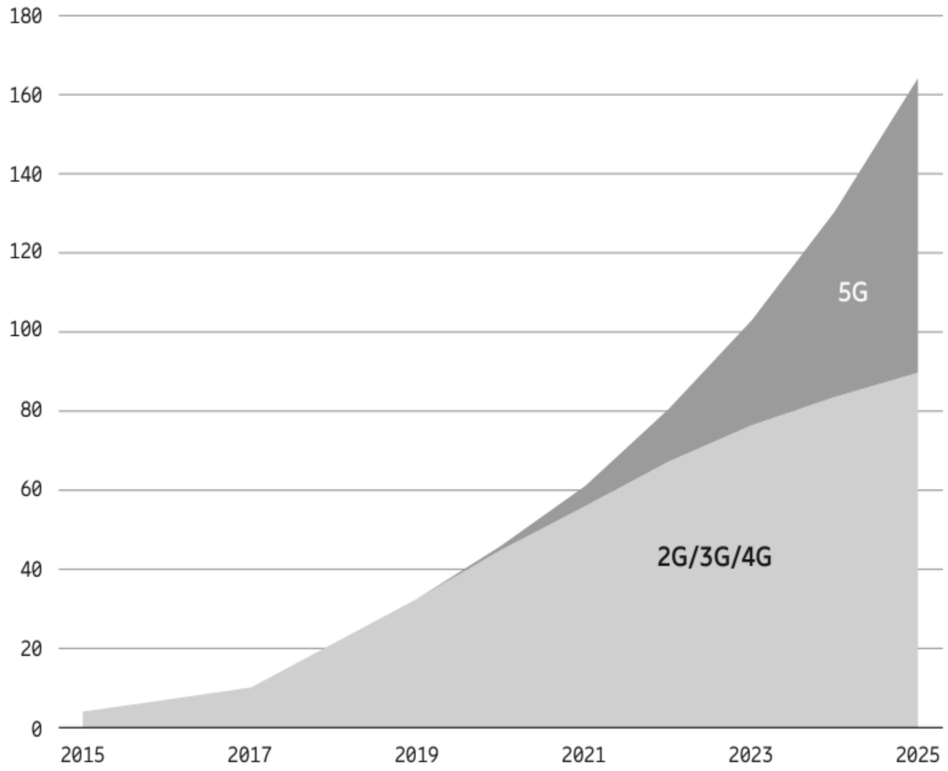

Figure 2.6: Global mobile data traffic (exabytes per month)

Smartphones continue to be at the epicenter of this development as they generate most of the mobile data traffic - about 95 percent, today, a share that is projected to increase throughout the forecast period. Populous markets that launch 5G early are likely to lead traffic growth over the forecast period. By 2025, we expect that 45 percent of total mobile data traffic will be carried by 5G networks.

Large variations in traffic growth across regions

Traffic growth can be very volatile between years, and can also vary significantly between countries, depending on local market dynamics. In the US, the traffic growth rate declined slightly during 2018 but recovered to previously expected rates during 2019. In China, 2018 was a year of record traffic growth. India's traffic growth continued its upward trajectory, and it remains the region with the highest usage per smartphone and per month. Globally, the growth in mobile data traffic per smartphone can be attributed to three main drivers: improved device capabilities, an increase in data-intensive content and more affordable data plans.

Around 410 million additional smartphone users are expected in India by 2025. In the India region, the average monthly mobile data usage per smartphone continues to show robust growth, boosted by the rapid adoption of 4G. Low prices for mobile broadband services, affordable smartphones and people's changing video viewing habits have continued to drive monthly usage growth in the region. According to GlobalData, India Telecom Operators Country Intelligence Report (2019), only 4 percent of households have fixed broadband, making smartphones the only way to access the internet in many cases.

Total traffic is projected to triple, reaching 21 EB per month in 2025. This comes from two factors: high growth in the number of smartphone users, including growth in rural areas, and an increase in average usage per smartphone. A total of around 410 million additional smartphone users are expected in India by 2025. Even if the traffic per existing smartphone user continues to grow significantly over time, the increase in average traffic per smartphone is expected to moderate as more consumers in India acquire smartphones. The average traffic per smartphone is expected to increase to around 25 GB per month in 2025.

2.4. REQUIREMENTS: QOS AND QOE

The notion of Quality of Service has served as a central research topic in communication networks for more than a decade, however, usually starting from a rather technical view on service quality. Therefore, the notion of Quality of Experience has emerged, redirecting the focus towards the end-user, and trying to quantify user's subjective experience gained from using a service.

2.4.1 Quality of Service

For at least a decade, Quality of Service (QoS) has been one of the dominating research topics in communication networks. Whereas the Internet originally has been conceived as a best-effort network, the introduction of QoS architectures like Integrated Services or

Differentiated Services was supposed to pave the way for high-quality real-time services like Voice-over-IP or video streaming and thus to increase the competitiveness of packet-based TCP/IP networks.

Technology-centered approach

The technology-centered approach mainly emphasizes the concept of QoS and has its strongest reference from the ITU (International Telecommunications Union).

The ITU Recommendation E.800 [i.31] is the key reference and states that QoS is the: *"Totality of characteristics of a telecommunications service that bear on its ability to satisfy stated and implied needs of the user of the service."*.

Although the ITU definition refers to user satisfaction, QoS is mainly used by technicians to define technical parameters of telecommunication applications such as network delay and packet loss. In addition, a focus on user satisfaction is rather limited because it is only one of many measures of user behavior with a communication service. For example, other measures include the time taken to perform a communication task (a measure of efficiency) and the accuracy with which a task is completed (a measure of effectiveness).

Commonly specified properties of QoS include:

- **Throughput**: A minimum or average throughput, in bytes per second or bits per second, for a given logical connection or traffic flow.
- **Delay**: The average or maximum delay. Also called latency.
- **Packet jitter**: Typically, the maximum allowable jitter.
- **Error rate**: Typically, maximum error rate, in terms of fraction of bits delivered in error.
- **Packet loss**: Fraction of packets lost.
- **Priority**: A network may offer a given number of levels of priority. The assigned level for various traffic flows influences the way in which the different flows are handled by the network.
- **Availability**: Expressed as a percentage of time available.
- **Security**: Different levels or types of security may be defined.

2.4.2 Quality of Experience

There are different definitions of Quality of Experience across current ITU, ETSI and other literature. A research document by ETSI defined Quality of Experience (QoE) as: *"A measure of user performance based on both objective and subjective psychological measures of using an ICT service or product."*

NOTE 1: It considers technical parameters (e.g. QoS) and usage context variables (e.g. communication task) and measures both the process and outcomes of communication (e.g. user effectiveness, efficiency, satisfaction, and enjoyment).

NOTE 2: The appropriate psychological measures will be dependent on the communication context.

Objective psychological measures do not rely on the opinion of the user (e.g. task completion time measured in seconds, task accuracy measured in number of errors). Subjective psychological measures are based on the opinion of the user (e.g. perceived quality of medium, satisfaction with a service). For example, a service provider may conclude that a service with a certain level of QoS used for a particular communication situation offers users excellent QoE, whist with a different level of QoS provides poor QoE.

For practical application, these features need to be converted to quantitative measures. The management of QoE has become a crucial concept in the deployment of future successful applications, services, and products.

The greatest challenges in providing QoE are developing effective methods for converting QoE features to quantitative measures and translating QoE measures to QoS measures. Whereas QoS can now easily be measured, monitored, and controlled at both the networking and application layers, and at both the end system and network sides, QoE is something that is still quite intricate to manage.

2.5 ROUTING

Routing and congestion control are the basic tools needed to support network traffic and to provide QoS and QoE. Both mechanisms are fundamental to the operation of a network and its capability to transmit and deliver packet traffic.

2.5.1 Characteristics

Routers forward packets from the original source to the destination. A router is considered a Layer 3 device because its primary forwarding decision is based on the information in the Layer 3 IP packet, specifically the destination IP address. This is known as routing and the decision is generally based on some performance criterion with the simplest one being minimum-hop route through the network.

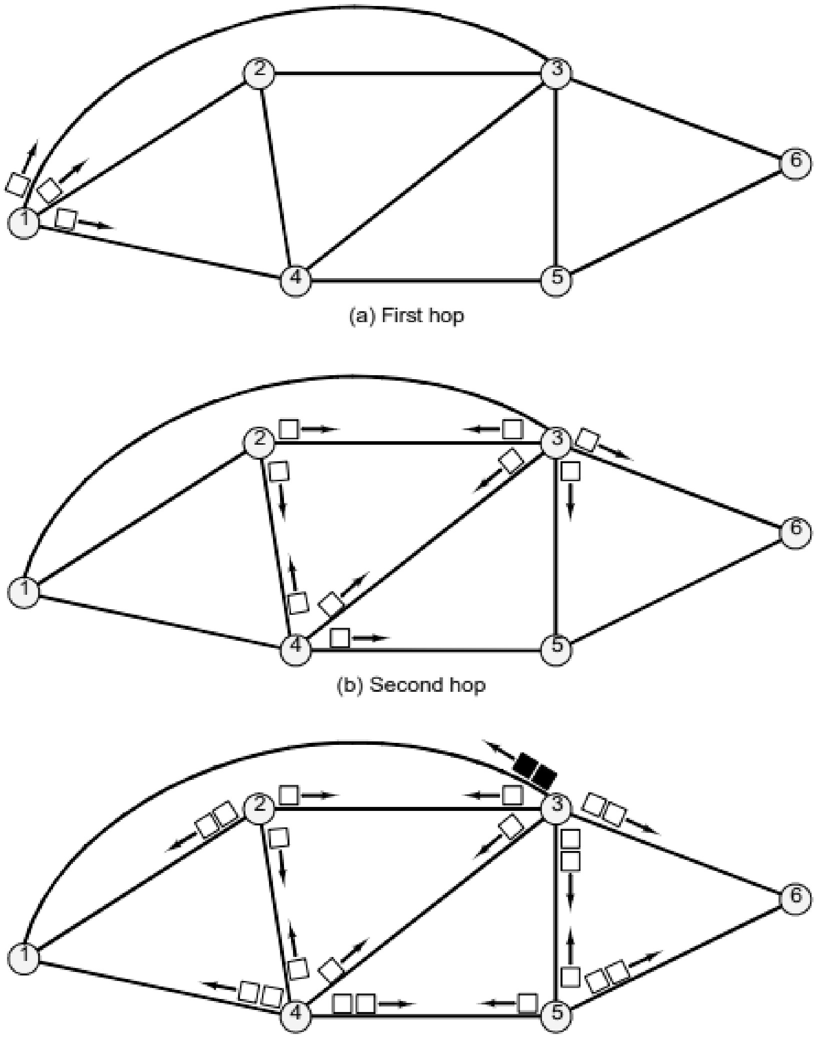

(a) First hop

(b) Second hop

(c) Third hop

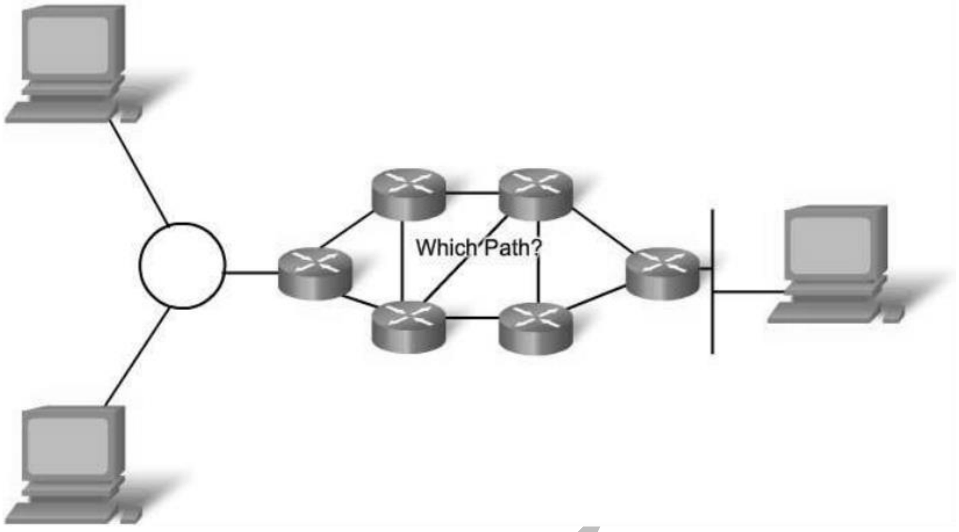

(d) Hypothetical Network Architecture

Figure 2.7: Network Architecture Example

A generalization of the minimum-hop criterion is least-cost routing. In this case, a cost is associated with each link, and, for any pair of attached stations, the route through the network that accumulates the least cost is sought. Figure 2.7 illustrates a network with numbers circled are nodes and lines connecting them represent links between these nodes. The shortest path from node 1 to node 6 is node 1 to node 3 to node 6 (1-3-6).

2.5.2 Packet Forwarding

The key function of any router is to accept incoming packets and forward them. For this purpose, a router maintains forwarding tables. A router's forwarding table shows, for each destination, the identity of the next node on the router. Each router may be responsible for discovering the appropriate routes. Alternatively, a network control center may be responsible for designing routes for all routers and maintaining a central forwarding table, providing each router with individual forwarding tables relevant only to that router.

Additional information is often used to determine the forwarding decision, such as the source address, packet flow identifier, or security level of the packet:

- **Failure**: When a node or link fails, it can no longer be used as part of a route.
- **Congestion**: When a particular portion of the network is heavily congested, it is desirable to route packets around rather than through the area of congestion.
- **Topology change**: The insertion of new links or nodes affects routing.

For adaptive routing to be possible, information about the state of the network must be exchanged among the nodes or between the nodes and a central controller.

2.6 CONGESTION CONTROL

Congestion occurs when the number of packets being transmitted through the network approaches the packet handling capacity of the network. Congestion control aims to keep number of packets below level at which performance falls off dramatically. It basically is reduced quality of service occurring when a network node or link is carrying more data than it can handle.

2.6.1 Effects of Congestion

Typical effects of congestion include queueing delay, packet loss or the blocking of new connections.

Queueing Delay:

In telecommunication and computer engineering, the queuing delay or queueing delay is the time a job waits in a queue until it can be executed. In a switched network, queuing delay is the time between the completion of signaling by the call originator and the arrival of a ringing signal at the call receiver. Queuing delay may be caused by delays at the originating switch, intermediate switches, or the call receiver servicing switch. In a data network, queuing delay is the sum of the delays between the request for service and the establishment of a circuit to the called data terminal equipment (DTE). In a packet-switched network, queuing delay is the sum of the delays encountered by a packet between the time of insertion into the network and the time of delivery to the address.

This term is most often used about routers. When packets arrive at a router, they must be processed and transmitted. A router can only process one packet at a time. If packets arrive faster than the router can process them (such as in a burst transmission) the router puts them into the queue (also called the buffer) until it can get around to transmitting them. Delay can also vary from packet to packet, so averages and statistics are usually generated when measuring and evaluating queuing delay.

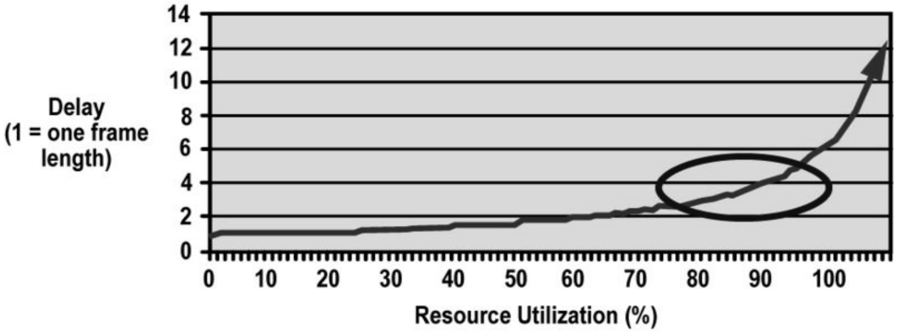

Figure 2.8: queuing delay based on throughput

Figure 2.8 displays the well-known effect of queuing delay based on throughput. As a queue begins to fill up due to traffic arriving faster than it can be processed, the amount of delay a particular packet experiences traversing the queue increases. The speed at which the contents of a queue can be processed is a function of the transmission rate of the facility. This leads to the classic "delay curve" depicted in the image to the right.

Packet Loss:

Packet loss occurs when one or more packets of data travelling across a computer network fail to reach their destination. Packet loss is either caused by errors in data transmission, typically across wireless networks, or network congestion. Packet loss is measured as a percentage of packets lost with respect to packets sent. In real-time applications like streaming media or online game, packet loss can affect a user's quality of experience (QoE). High packet loss rate indicates that users sustain undoubtedly a very poor quality.

Packet loss is also closely associated with quality of service (QoS) considerations. The amount of packet loss that is acceptable depends on the type of data being sent. A typical service level agreement (SLA) between Service Provider and Recipient (Company) could also mention a clause on Network packet delivery (reliability) as Average monthly packet loss shall be no greater than 0.1 percent (or successful delivery of 99.9 percent of packets). In such cases, packet loss could be defined as the percentage of packets that are dropped between backbone hubs on the Service Provider's Network.

Packet loss is detected by reliable protocols such as TCP. Reliable protocols react to packet loss automatically, so when a person such as a network administrator needs to detect and diagnose packet loss, they typically use status information from network equipment or purpose-built tools. The Internet Control Message Protocol provides an echo functionality, where a special packet is transmitted that always produces a reply. Tools such as ping, traceroute, and MTR use this protocol to provide a visual representation of the path packets are taking, and to measure packet loss at each hop. Many routers have status pages or logs, where the owner can find the number or percentage of packets dropped over a particular period.

Blocking of new connections:

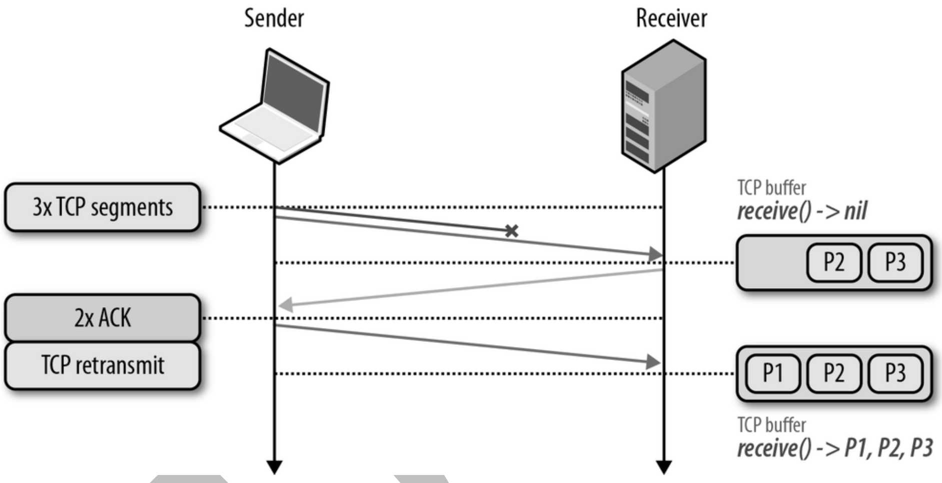

Figure 2.9: TCP Head-of-line blocking

2.6.2 Congestion Control Techniques

Congestion control techniques can be broadly classified into two categories: open loop congestion control and closed loop congestion control.

A. Open loop congestion control:

Open loop congestion control policies are applied to prevent congestion before it happens. The congestion control is handled either by the source or the destination.

Policies adopted by open loop congestion control:

1. Retransmission Policy:
   It is the policy in which retransmission of the packets are taken care. If the sender feels that a sent packet is lost or corrupted, the packet needs to be retransmitted. This transmission may increase the congestion in the network. To

prevent congestion, retransmission timers must be designed to prevent congestion as also able to optimize efficiency.

2. Window Policy:
   The type of window at the sender side may also affect the congestion. Several packets in the Go-back-n window are resent, although some packets may be received successfully at the receiver side. This duplication may increase the congestion in the network and making it worse. Therefore, Selective repeat window should be adopted as it sends the specific packet that may have been lost.

3. Discarding Policy:
   A good discarding policy adopted by the routers is that the routers may prevent congestion and at the same time partially discards the corrupted or less sensitive package as also able to maintain the quality of a message. In case of audio file transmission, routers can discard less-sensitive packets to prevent congestion as also maintain the quality of the audio file.

4. Acknowledgment Policy:
   Since acknowledgement are also the part of the load in network, the acknowledgment policy imposed by the receiver may also affect congestion. Several approaches can be used to prevent congestion related to acknowledgment. The receiver should send acknowledgement for N packets rather than sending acknowledgement for a single packet. The receiver should send an acknowledgment only if it must send a packet or a timer expires.

5. Admission Policy:
   In admission policy a mechanism should be used to prevent congestion. Switches in a flow should first check the resource requirement of a network flow before transmitting it further. If there is a chance of a congestion or is a congestion in the network, router should deny establishing a virtual network connection to prevent further congestion.

All the above policies are adopted to prevent congestion before it happens in the network.

B. Closed Loop Congestion Control

Following closed loop congestion control techniques are used to treat or alleviate congestion after it happens.

1. Backpressure:

Backpressure is a technique in which a congested node stop receiving packet from upstream node. This may cause the upstream node or nodes to become congested and rejects receiving data from above nodes. Backpressure is a node-to-node congestion control technique that propagate in the opposite direction of data flow. The backpressure technique can be applied only to virtual circuit where each node has information of its above upstream node.

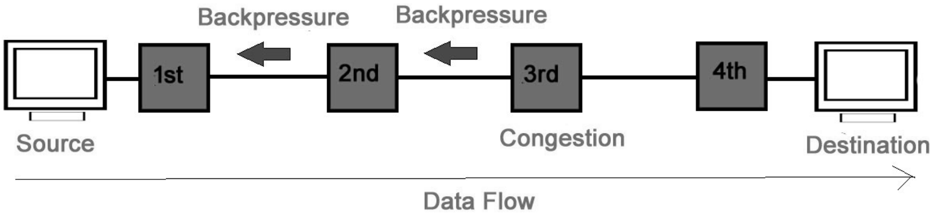

Figure 2.10 Backpressure

In figure 2.10, third node is congested and stops receiving packets as a result second node may be get congested due to slowing down of the output data flow. Similarly, first node may get congested and informs the source to slow down.

2. Choke Packet Technique:
   Choke packet technique is applicable to both virtual networks as well as datagram subnets. A choke packet is a packet sent by a node to the source to inform it of congestion. Each router monitors its resources and the utilization at each of its output lines. whenever the resource utilization exceeds the threshold value, which is set by the administrator, the router directly sends a choke packet to the source giving it a feedback to reduce the traffic. The intermediate nodes through which the packet has travelled are not warned about congestion.

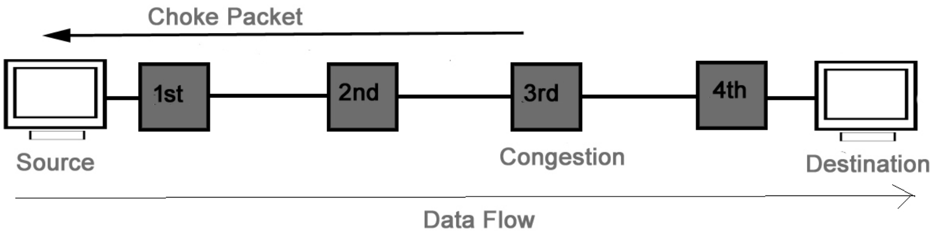

Figure 2.11 Choke Packet

3. Implicit Signaling:
   In implicit signaling, there is no communication between the congested nodes and the source. The source guesses that

there is congestion in a network. For example, when sender sends several packets and there is no acknowledgment for a while, one assumption is that there is a congestion.

4. Explicit Signaling:
   In explicit signaling, if a node experiences congestion it can explicitly sends a packet to the source or destination to inform about congestion. The difference between choke packet and explicit signaling is that the signal is included in the packets that carry data rather than creating different packet as in case of choke packet technique.

   Explicit signaling can occur in either forward or backward direction.

   - Forward Signaling:
     A bit can be set in a packet moving in the direction of the congestion. This bit can warn the destination that there is congestion. The receiver in this case can use policies, such as slowing down the acknowledgments, to alleviate the congestion.

   - Backward Signaling:
     A bit can be set in a packet moving in the direction opposite to the congestion. This bit can warn the source that there is congestion and that it needs to slow down to avoid the discarding of packets.

   Explicit congestion signaling approaches could be divided into three general categories:

   A. Binary: A bit is set in a data packet as it is forwarded by the congested node. When a source receives a binary indication of congestion on a logical connection, it may reduce its traffic flow.

   B. Credit based: These schemes are based on providing an explicit credit to a source over a logical connection. The credit indicates how many octets or how many packets the source may transmit. When the credit is exhausted, the source must await additional credit before sending additional data. Credit-based schemes are common for end-to-end flow control, in which a destination system uses credit to prevent the source from overflowing the destination buffers, but credit-based schemes have also been considered for congestion control. Credit-based schemes are defined in Frame Relay and ATM networks.

C. Rate based: These schemes are based on providing an explicit data rate limit to the source over a logical connection. The source may transmit data at a rate up to the set limit. To control congestion, any node along the path of the connection can reduce the data rate limit in a control message to the source.

2.7 SDN AND NFV

The Internet is the midst of a transformation, one that moves away from bundled proprietary devices, and instead embraces disaggregating network hardware (which becomes commodity) from the software that controls it (which scales in the cloud). The transformation is generally known as Software-Defined Networking (SDN).

2.7.1 Software-Defined Networking

Software-defined networks provide an enhanced level of flexibility and customizability to meet the needs of newer networking and IT trends such as cloud, mobility, social networking, and video.

The SDN Architecture is:

- **Directly programmable**
  Network control is directly programmable because it is decoupled from forwarding functions.

- **Agile**
  Abstracting control from forwarding lets administrators dynamically adjust network-wide traffic flow to meet changing needs.

- **Centrally managed**
  Network intelligence is (logically) centralized in software based SDN controllers that maintain a global view of the network, which appears to applications and policy engines as a single, logical switch.

- **Programmatically configured**
  SDN lets network managers configure, manage, secure, and optimize network resources very quickly via dynamic, automated SDN programs, which they can write themselves because the programs do not depend on proprietary software.

- **Open standards-based and vendor-neutral**

When implemented through open standards, SDN simplifies network design and operation because instructions are provided by SDN controllers instead of multiple, vendor-specific devices and protocols.

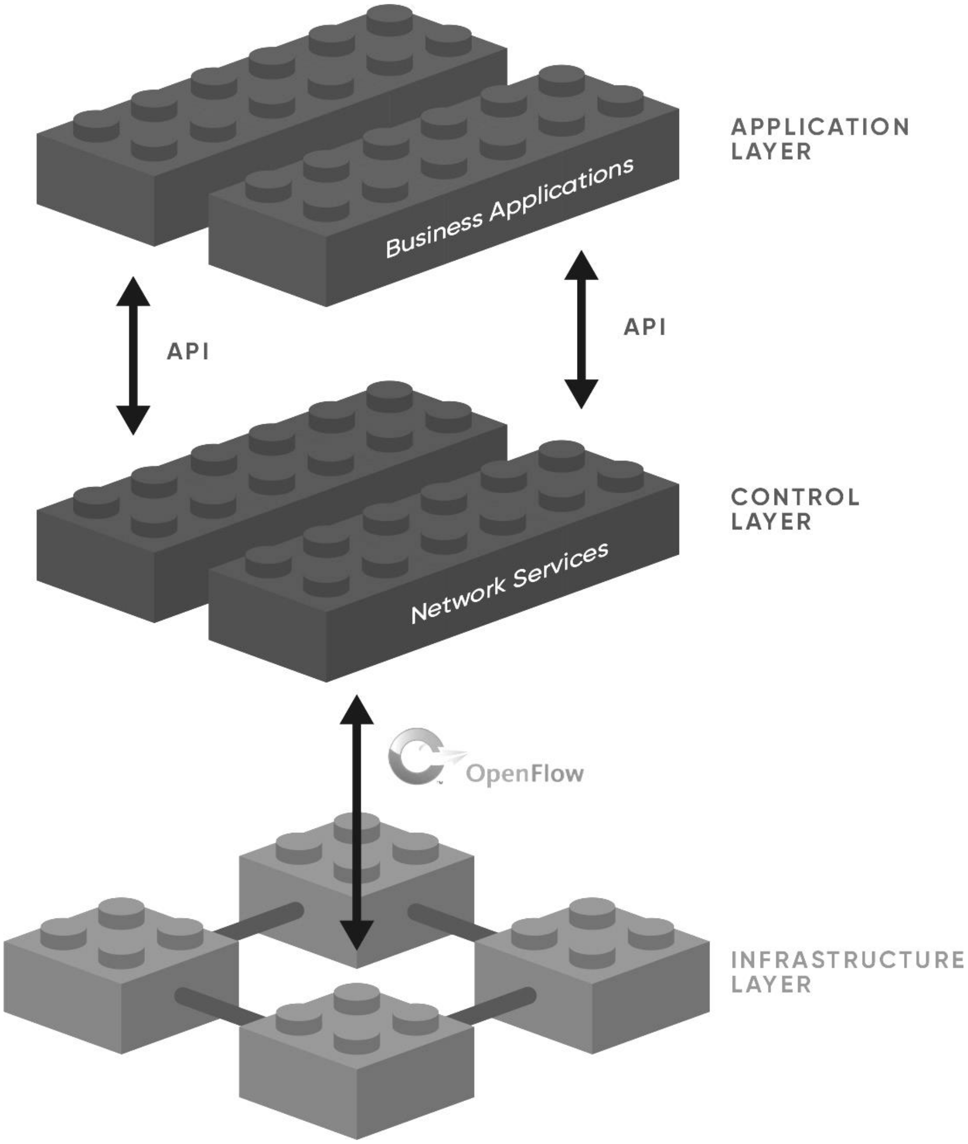

Figure 2.12 SDN Architecture

SDN Functionality

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, and so on). Control in such a traditional network is exercised by means of a routing and control network protocol that is implemented in each network node. This approach is relatively inflexible and requires all 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.

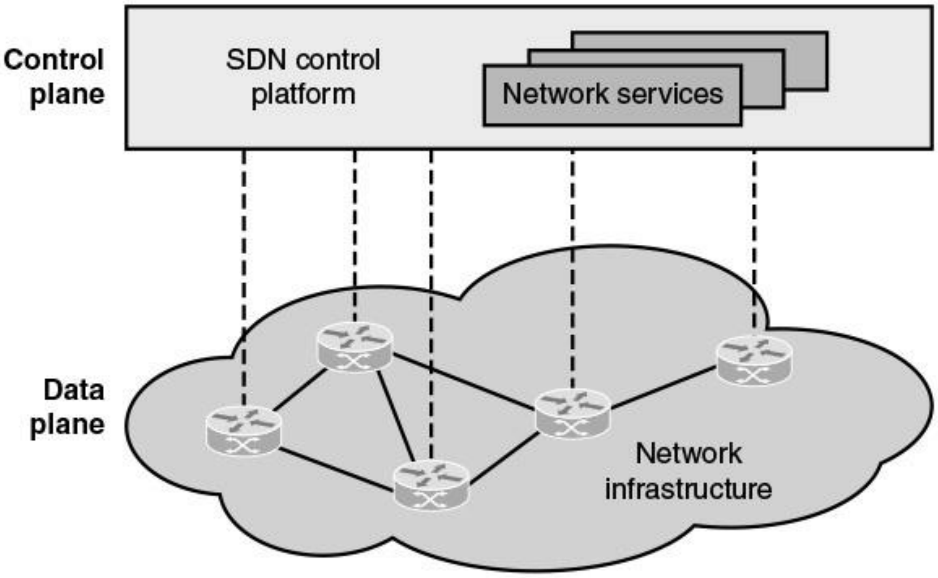

Figure 2.13 Software-Defined Networking

This 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 function subsumed by the controller, switches simply manage flow tables whose entries can only be populated by the controller. The switches constitute the data plane.
Communication between the controller and the switches uses a standardized protocol.

2.7.2 Network Functions Virtualization

Virtual machine technology over the Internet or an enterprise network has been used for application-level server functions such as database servers, cloud servers, web servers, e-mail servers, and so on. This same technology, however, can equally be applied to network devices, such as routers, LAN switches, firewalls, and IDS/IPS servers.

Network Functions Virtualization (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 utilizes 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. NFV has several features in common with SDN.

They share the following objectives:

- Move functionality to software
- Use commodity hardware platforms instead of proprietary platforms
- Use standardized or open application program interfaces (APIs)
- Support more efficient evolution, deployment, and repositioning of network functions

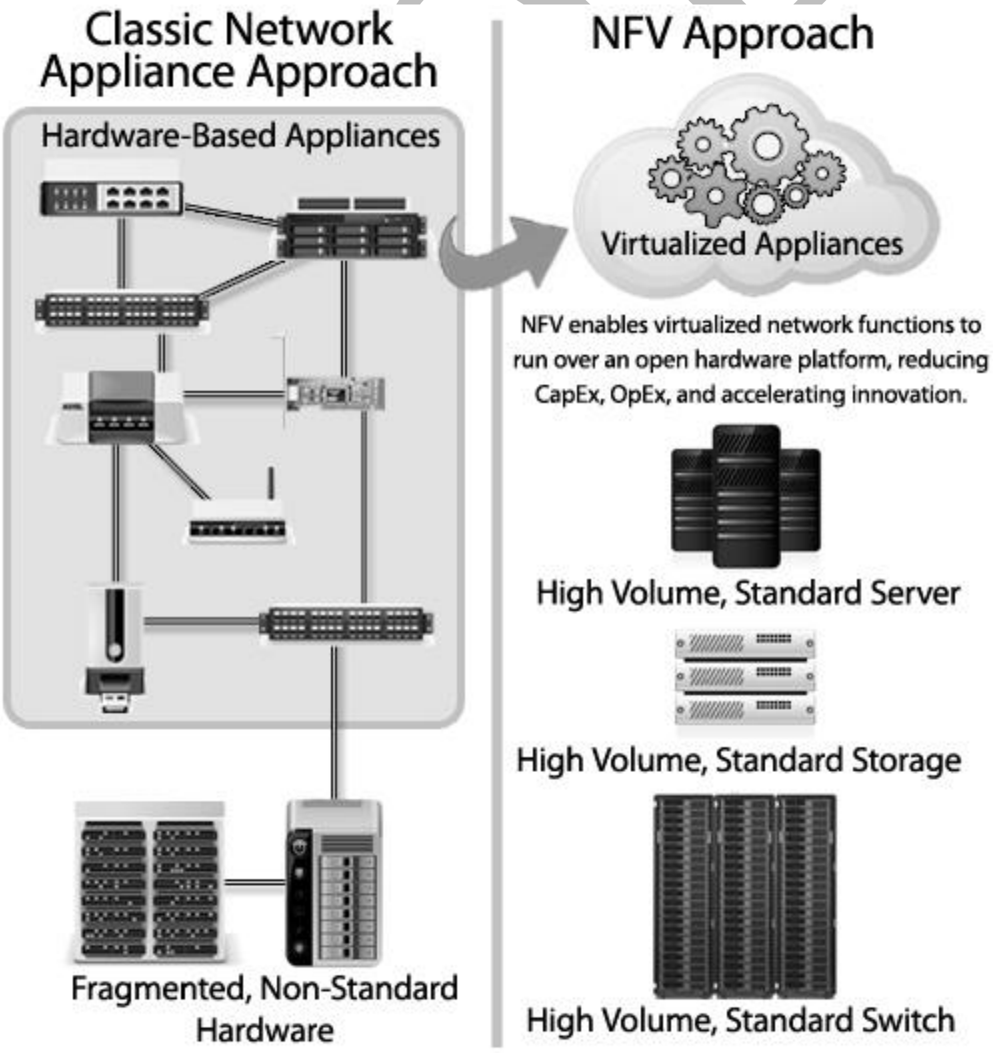
Classic Network Appliance Approach
Hardware-Based Appliances
Fragmented, Non-Standard Hardware
NFV Approach
Virtualized Appliances
NFV enables virtualized network functions to run over an open hardware platform, reducing CapEx, OpEx, and accelerating innovation.
High Volume, Standard Server
High Volume, Standard Storage
High Volume, Standard Switch

Figure 2.14 Network Functions Virtualization approach

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. So, either can be used alone, but the two can be combined to reap greater benefits.

NFV reduces the need for dedicated hardware to deploy and manage networks by offloading network functions into software that can run on industry-standard hardware and can be managed from anywhere within the operator's network.

Separating network functions from hardware yields numerous benefits for the network operator, which include:

- Reduced space needed for network hardware
- Reduce network power consumption
- Reduced network maintenance costs
- Easier network upgrades
- Longer life cycles for network hardware
- Reduced maintenance and hardware costs

2.8 MODERN NETWORKING ELEMENTS

Ultimately, the concern of a network service provider is about the set of network devices (such as routers) and the control and management of the functions they perform (such as packet forwarding). If NFV is used, these network functions are implemented in software and executed on VMs. If instead the network functions are implemented on dedicated machines and SDN is used, the control functions are implemented on central SDN controllers, which interact with the network devices. However, SDN and NFV are not mutually exclusive.

If both SDN and NFV are implemented for a network, the following relationships hold:

- Network data plane functionality is implemented on VMs.
- The control plane functionality may be implemented on a dedicated SDN platform or on an SDN VM.

In either case, the SDN controller interacts with the data plane functions running on VMs.

QoS measures are commonly used to specify the service required by various network customers or users and to dictate the traffic management policies used on the network. The common case, until recently, is that QoS was implemented on network that used neither NFV nor SDN. In this case, routing and traffic control policies must be configured directly on network devices using a variety of automated and manual techniques. If NFV but not SDN is implemented, the QoS settings are communicated to the VMs. With SDN, regardless of whether NFV is used, it is the SDN controller that is responsible for enforcing QoS parameters for the various network users. If QoE considerations come into play, these are used to adjust QoS parameters to satisfy the users' QoE requirements.

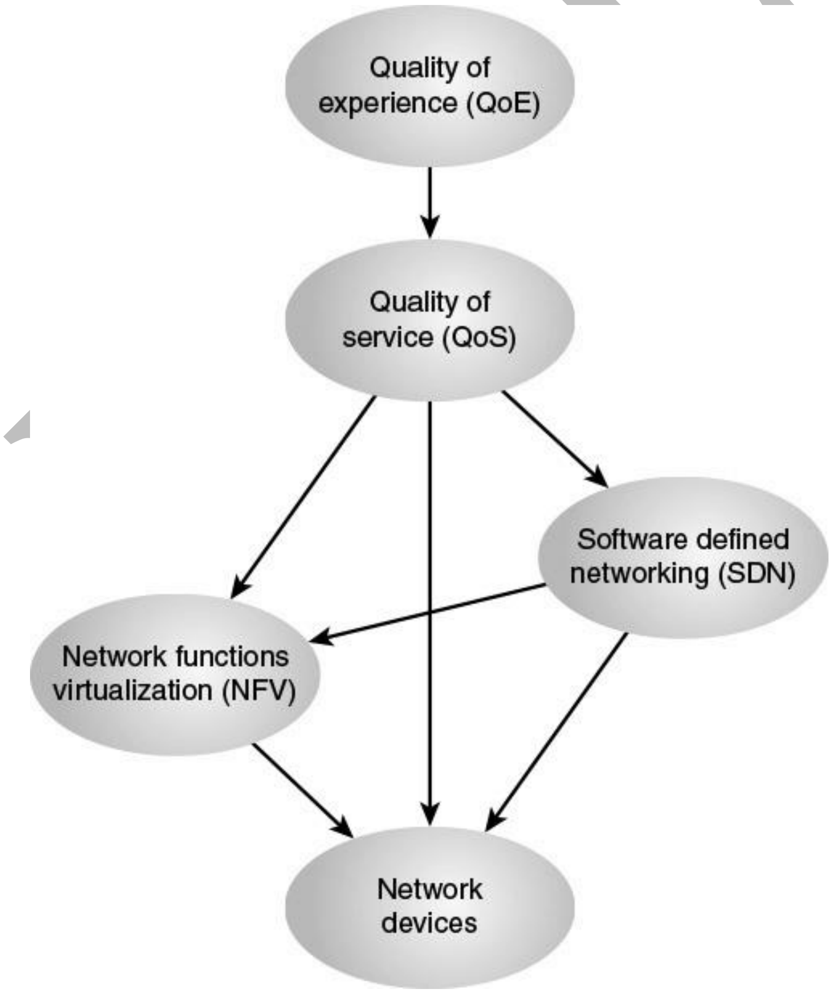
Quality of experience (QoE)
Quality of service (QoS)
Software defined networking (SDN)
Network functions virtualization (NFV)
Network devices

Figure 2.15 Modern Networking Schema

2.9 SUMMARY

- Elastic traffic is that which can adjust, over wide ranges, to changes in delay and throughput across an internet and still meet the needs of its applications
- Inelastic traffic does not easily adapt, if at all, to changes in delay and throughput across an internet
- big data refers to everything that enables an organization to create, manipulate, and manage very large data sets (measured in terabytes, petabytes, exabytes, and so on) and the facilities in which these are stored.
- Traditional business data storage and management technologies include relational database management systems (RDBMS), network-attached storage (NAS), storage-area networks (SANs), data warehouses (DWs), and business intelligence (BI) analytics.
- A cloud-based network is an enterprise network that can be extended to the cloud
- The technology-centered approach mainly emphasizes the concept of QoS and has its strongest reference from the ITU (International Telecommunications Union).
- Routing and congestion control are the basic tools needed to support network traffic and to provide QoS and QoE.
- Packet loss occurs when one or more packets of data travelling across a computer network fail to reach their destination
- Congestion control techniques can be broadly classified into two categories: open loop congestion control and closed loop congestion control
- Software-defined networks provide an enhanced level of flexibility and customizability to meet the needs of newer networking and IT trends such as cloud, mobility, social networking, and video.
- Network Functions Virtualization (NFV) decouples network functions, such as routing, firewalling, intrusion detection, and Network Address Translation from proprietary hardware platforms and implements these functions in software.

2.10 REVIEW QUESTION

1. Present an overview of the major categories of packet traffic on the Internet and internets, including elastic, inelastic, and real-time traffic.

2. Discuss the traffic demands placed on contemporary networks by big data, cloud computing, and mobile traffic.
3. Explain the concept of quality of service.

4. Explain the concept of quality of experience.

5. Understand the essential elements of routing.

6. Understand the effects of congestion and the types of techniques used for congestion control.

7. Compare and contrast software-defined networking and network functions virtualization.

8. What is congestion? Why does it occur?

9. What is choke packet? How is it used for congestion control?

2.11 REFERENCES

1. Foundations of Modern Networking: SDN, NFV, QoE, IoT, and Cloud, First printing: November 2015 by William Stallings, Copyright © 2016 by Pearson Education, Inc. ISBN-13: 978-0-13-417539-3

2. Research paper “On Traffic Types and Service Classes in the Internet” by Mansour J. Karam, Fouad A. Tobagi, available online at http://mmnetworks.stanford.edu/papers/karam_globecom00.pdf

3. Center for Applied Internet Data Analysis (CAIDA) research material on Internet Traffic Classification published and available online at https://www.caida.org/research/traffic-analysis/classification-overview/

4. Cloud Computing Networking Theory, Practice, and Development by Lee Chao, © 2016 by Taylor & Francis Group, LLC. ISBN-13: 978-1-4822-5482-2

5. Cloud Computing: A Practical Approach by Anthony T. Velte Toby J. Velte, Robert Elsenpeter, Copyright © 2010 by The McGraw-Hill Companies. ISBN-13: 978-0-07-162695-8

6. Mobile network traffic Q1 2020 and Mobile data traffic outlook sections adopted from Ericsson Mobility Report June 2020, Publisher: Fredrik Jejdling, available online at https://www.ericsson.com/49da93/assets/local/mobility-report/documents/2020/june2020-ericsson-mobility-report.pdf

7. Research paper "From Quality of Service to Quality of Experience" by Markus Fiedler, Kalevi Kilkki and Peter Reichl published online at https://drops.dagstuhl.de/opus/volltexte/2009/2235/pdf/09192.SWM.Paper.2235.pdf

8. Technical Report on "Human Factors (HF); Quality of Experience (QoE) requirements for real-time communication services" by ETSI, © European Telecommunications Standards Institute 2009, available online at https://www.etsi.org/deliver/etsi_tr/102600_102699/102643/01.00.01_60/tr_102643v010001p.pdf

9. Technical Article on Queuing delay, published by hill associates, archive available online at https://web.archive.org/web/20150904041151/http://www.hill2dot0.com/wiki/index.php?title=Queuing_delay

10. Technical Article on Queuing delay, published by Wikipedia available online at https://en.wikipedia.org/wiki/Queuing_delay

11. Technical Article on Packet Loss, published by Wikipedia available online at https://en.wikipedia.org/wiki/Packet_loss

12. Chapter on Building Blocks of TCP, published by O'Reilly at https://www.oreilly.com/library/view/high-performance-browser/9781449344757/ch02.html

13. Computer Network: Lecture Notes, prepared by Mr. Daya Ram Budhathoki, available online at https://dayaramb.files.wordpress.com/2011/03/computer-network-notes-pu.pdf

14. SDN Architecture https://opennetworking.org/sdn-definition/

15. Network Functions Virtualizations approach available online at: https://www.blueplanet.com/resources/What-is-NFV-prx.html

❖❖❖❖

Unit 4: Chapter 1

QUALITY OF SERVICE (QoS) AND USER QUALITY OF EXPERIENCE (QoE)

Unit Structure

1.0 OBJECTIVES:

- Describe the ITU-T QoS architectural framework.
- Summarize the key concepts of the Integrated Services Architecture.
- Compare and contrast elastic and inelastic traffic.
- Explain the concept of differentiated services.
- Understand the use of service level agreements.
- Describe IP performance metrics.
- Present an overview of OpenFlow QoS support
- Explain the motivations for QoE.
- Define QoE.
- Explain the factors that could influence QoE.
- Present an overview of how QoE can be measured, including a discussion of the differences between
- Subjective and objective assessment.
- Discuss the various application areas of QoE.

1.1 INTRODUCTION:

Fundamental to the acceptance and success of any complex shared networking architecture is that it meets users expectation for performance. Traditionally, the means of defining expected performance, measuring it, providing it, and entering into well-defined agreements relating to it has been the concept of quality of service (QoS).

The Internet and enterprise IP-based networks continue to see rapid growth in the volume and variety of data traffic. Cloud computing, big data, the pervasive use of mobile devices on enterprise networks, and the increasing use of video streaming all contribute to the increasing difficulty in maintaining satisfactory network performance. Two key tools in measuring the network performance that an enterprise desires to achieve are quality of service (QoS) and quality of experience (QoE).

QoS and QoE enable the network manager to determine whether the network is meeting user needs and to diagnose problem areas that require adjustment to network management and network traffic control. There is a strong need to be able to support a variety of traffic, with a variety of QoS requirements, on IP- based networks.

1.2 BACKGROUND

The Internet and other IP-based networks provided a **best effort** delivery service. This means that the network attempts to allocate its resources with equal availability and priority to all traffic flows, with no regard for application priorities, traffic patterns and load, and customer requirements. To protect the network from congestion collapse and to guarantee that some flows do not crowd out other flows, congestion control mechanisms were introduced, which tended to throttle traffic that consumed excessive resources.

One of the most important congestion control techniques, introduced early on and still in wide use, is the TCP congestion control mechanism. TCP congestion control has become increasingly complex and sophisticated, but it is worth briefly summarizing the principles involved here. For each TCP connection between two end systems across a network, in each direction, a concept known as sliding window is used. TCP segments on a connection are numbered sequentially. The sending and receiving TCP entities maintain a window, or buffer, that defines the range of sequence numbered segments that may be transmitted. As segments arrive and are processed by the receiver, the receiver returns an acknowledgment indicating which segments have been received and implicitly indicated to the sender that the window of sequence numbers has advanced to allow more segments to be sent. Various algorithms are used by the sender to deduce the amount of congestion on a connection based on the round-trip delay for acknowledgments plus whether an acknowledgment is even received for a particular segment. As congestion is detected, the sending TCP entity reduces its transmission of segments to help ease congestion on the intervening network.

Although TCP congestion control and other network congestion control techniques can reduce the risk of excessive congestion, these techniques do not directly address QoS requirements. As the intensity and variety of traffic increased, various QoS mechanisms were developed, including Integrated Services Architecture (ISA) and differentiated services (DiffServ), accompanied by service level agreements (SLAs) so that the service provided to various customers was tunable and somewhat predictable. These mechanisms and services serve two purposes:

- Allocate network resources efficiently so as to maximize effective capacity
- Enable networks to offer different levels of QoS to customers on the basis of customer requirements

In this more sophisticated environment, the term ***best effort refers*** not to the network service as a whole but to a class of traffic treated in best effort fashion. All packets in the best effort traffic class are transmitted with no guarantee regarding the speed with which the packets will be transmitted to the recipient or that the data will even be delivered entirely.

1.3 QoS ARCHITECTURAL FRAMEWORK

The Y.1291 framework consists of a set of generic network mechanisms for controlling the network service response to a service request, which can be specific to a network element, or for signaling between network elements, or for controlling and administering traffic across a network. Figure 1.1 shows the relationship among these elements, which are organized into three planes: data, control, and management.

1.3.1 Data Plane

The data plane includes those mechanisms that operate directly on flows of data.

- **Traffic classification** refers to the assignment of packets to a traffic class by the ingress router at the ingress edge of the network. Typically, the classification entity looks at multiple fields of a packet, such as source and destination address, application payload, and QoS markings, and determines the aggregate to which the packet belongs. The flow label in the IPv6 header can be used for traffic classification. Other routers en route perform a classification function as well, but the classification does not change as the packets traverse the network.

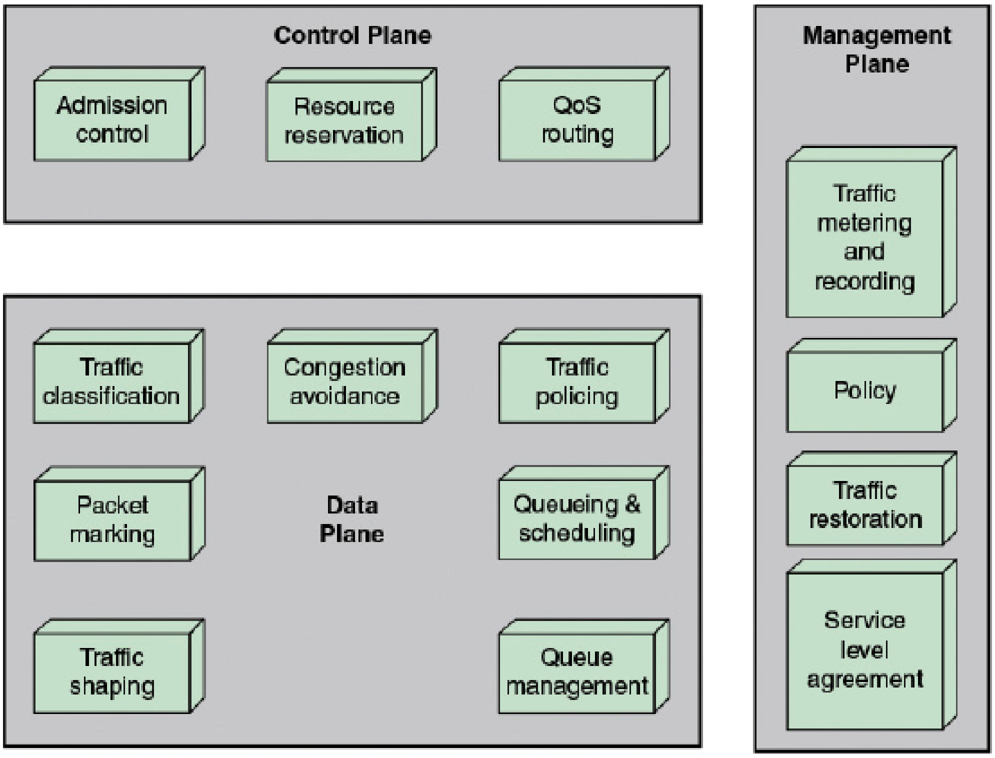

Figure 1.1 Architectural Framework for QoS Support

- **Packet marking** encompasses two distinct functions. First, packets may be marked by ingress edge nodes of a network to indicate some form of QoS that the packet should receive. An example is the Differentiated Services (DiffServ) field in the IPv4 and IPv6 packets and the Traffic Class field in MPLS labels. An ingress edge node can set the values in these fields to indicate a desired QoS. Such markings may be used by intermediate nodes to provide differential treatment to incoming packets. Second, packet marking can also be used to mark packets as nonconformant, either by the ingress node or intermediate nodes, which may be dropped later if congestion is experienced.
- **Traffic shaping** controls the rate and volume of traffic entering and transiting the network on a per-flow basis. The entity responsible for traffic shaping buffers nonconformant packets until it brings the respective aggregate in compliance with the traffic limitations for this flow.
- **Congestion avoidance** deals with means for keeping the load of the network under its capacity such that it can operate at an acceptable performance level. The specific objectives are to avoid significant queuing delays and, especially, to avoid congestion collapse. A typical congestion avoidance scheme acts by senders reducing the amount of traffic entering the network upon an indication that network congestion is occurring (or about to occur).
- **Traffic policing** determines whether the traffic being presented is, on a hop-by-hop basis, compliant with prenegotiated policies or contracts. Nonconformant packets may be dropped, delayed, or labeled as nonconformant.
- **Queuing and scheduling** algorithms, also referred to as queuing discipline algorithms, determine which packet to send next and are used primarily to manage the allocation of transmission capacity among flows.

- **Queue management** algorithms manage the length of packet queues by dropping packets when necessary or appropriate. Active management of queues is concerned primarily with congestion avoidance. One noteworthy example of queue management is random early detection (RED). RED drops incoming packets probabilistically based on an estimated average queue size. The probability for dropping increases as the estimated average queue size grows. There are a number of variants of RED that are in more common use than the original RED, with weighted RED (WRED) perhaps the most commonly implemented.

1.3.2 Control Plane

The control plane is concerned with creating and managing the pathways through which user data flows. It includes admission control, QoS routing, and resource reservation.

- **Admission control** determines what user traffic may enter the network. This may be in part determined by the QoS requirements of a data flow compared to the current resource commitment within the network. But beyond balancing QoS requests with available capacity to determine whether to accept a request, there are other considerations in admission control.
- **QoS routing** determines a network path that is likely to accommodate the requested QoS of a flow. This contrasts with the philosophy of the traditional routing protocols, which generally are looking for a least-cost path through the network.
- **Resource reservation** is a mechanism that reserves network resources on demand for delivering desired network performance to a requesting flow. An example of a protocol that uses this capability is the Resource Reservation Protocol (RSVP).

1.3.3 Management Plane

The management plane contains mechanisms that affect both control plane and data plane mechanisms. The control plane deals with the operation, administration, and management aspects of the network. It includes SLAs, traffic restoration, traffic metering and recording, and policy.

A **service level agreement (SLA)** typically represents the agreement between a customer and a provider of a service that specifies the level of availability, serviceability, performance, operation, or other attributes of the service.

- **Traffic metering and recording** concerns monitoring the dynamic properties of a traffic stream using performance metrics such as data rate and packet loss rate. It involves observing traffic characteristics at a given network point and collecting and storing the traffic information for analysis and further action. Depending on the conformance level, a meter can invoke necessary treatment (for example, dropping or shaping) for the packet stream.
- **Traffic restoration** refers to the network response to failures. This encompasses a number of protocol layers and techniques.
- **Policy** is a category that refers to a set of rules for administering, managing, and controlling access to network resources. They can be specific to the needs of the service provider or reflect the agreement between the customer and service provider, which may include reliability and availability requirements over a period of time and other QoS requirements.

1.4 INTEGRATED SERVICES ARCHITECTURE (ISA)

1.4.1 ISA Approach

The purpose of ISA is to enable the provision of QoS support over IP-based internets. The central design issue for ISA is how to share the available capacity in times of congestion. For an IP-based Internet that provides only a best effort service, the tools for controlling congestion and providing service are limited. In essence, routers have two mechanisms to work with:

- **Routing algorithm**: Some routing protocols in use in internets allow routes to be selected to minimize delay. Routers exchange information to get a picture of the delays throughout the Internet. Minimum-delay routing helps to balance loads, thus decreasing local congestion, and helps to reduce delays seen by individual TCP connections.
- **Packet discard**: When a router’s buffer overflows, it discards packets. Typically, the most recent packet is discarded. The effect of lost packets on a TCP connection is that the sending TCP entity backs off and reduces its load, thus helping to alleviate Internet congestion.

In ISA, each IP packet can be associated with a flow. A flow is a distinguishable stream of related IP packets that results from a single user activity and requires the same QoS. ISA makes use of the following functions to manage congestion and provide QoS transport:

- **Admission control**: For QoS transport (other than default best effort transport), ISA requires that a reservation be made for a new flow. If the routers collectively determine that there are insufficient resources to guarantee the requested QoS, the flow is not admitted. The protocol RSVP is used to make reservations.
- **Routing algorithm**: The routing decision may be based on a variety of QoS parameters, not just minimum delay.
- **Queuing discipline**: A vital element of the ISA is an effective queuing policy that takes into account the differing requirements of different flows.
- **Discard policy**: A discard policy determines which packets to drop when a buffer is full and new packets arrive.

1.4.2 ISA Components

Figure 1.2 is a general depiction of the implementation architecture for ISA within a router. Below the thick horizontal line are the forwarding functions of the router; these are executed for each packet and therefore must be highly optimized. The remaining functions, above the line, are background functions that create data structures used by the forwarding functions.

The principal background functions are as follows:

- **Reservation protocol**: This protocol reserves resources for a new flow at a given level of QoS. It is used among routers and between routers and end systems. The reservation protocol is responsible for maintaining flow-specific state information at the end systems and at the routers along the path of the flow. RSVP is used for this purpose.

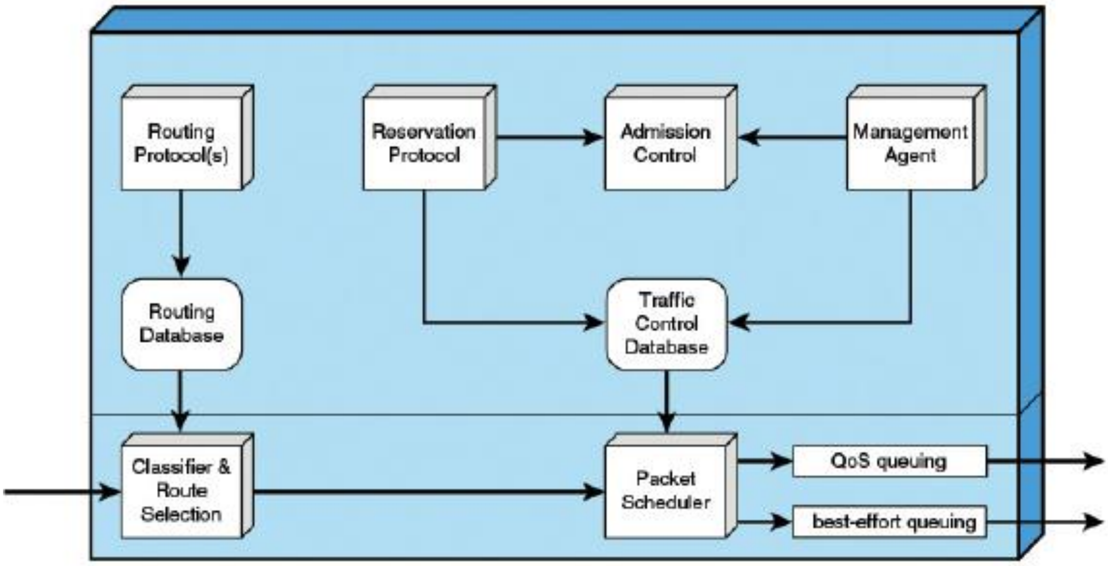

Figure 1.2 Integrated Services Architecture Implemented in Router

- **Admission control**: When a new flow is requested, the reservation protocol invokes the admission control function. This function determines if sufficient resources are available for this flow at the requested QoS. This determination is based on the current level of commitment to other reservations or on the current load on the network.
- **Management agent**: A network management agent can modify the traffic control database and to direct the admission control module to set admission control policies.
- **Routing protocol**: The routing protocol is responsible for maintaining a routing database that gives the next hop to be taken for each destination address and each flow.

These background functions support the main task of the router, which is the forwarding of packets. The two principal functional areas that accomplish forwarding are the following:

1. **Classifier and route selection**: For the purposes of forwarding and traffic control, incoming packets must be mapped into classes. A class may correspond to a single flow or to a set of flows with the same QoS requirements. For example, the packets of all video flows or the packets of all flows attributable to a particular organization may be treated identically for purposes of resource allocation and queuing discipline. The selection of class is based on fields in the IP header. Based on the packet's class and its destination IP address, this function determines the next-hop address for this packet.
2. **Packet scheduler**: This function manages one or more queues for each output port. It determines the order in which queued packets are transmitted and the selection of packets for discard, if necessary. Decisions are made based on a packet's class, the contents of the traffic control database, and current and past activity on this outgoing port. Part of the packet scheduler's task is that of policing, which is the function of determining whether the packet traffic in a given flow exceeds the requested capacity and, if so, deciding how to treat the excess packets.

1.4.3 ISA Services

ISA service for a flow of packets is defined on two levels. First, a number of general categories of service are provided, each of which provides a certain general type of service guarantees. Second, within each category, the service for a particular flow is specified by the values of certain parameters; together, these values are referred to as a traffic specification (TSpec). Three categories of service are defined:

1. Guaranteed
2. Controlled load
3. Best effort

An application can request a reservation for a flow for a guaranteed or controlled load QoS, with a TSpec that defines the exact amount of service required. If the reservation is accepted, the TSpec is part of the contract between the data flow and the service.

Guaranteed Service

The key elements of the guaranteed service are as follows:

- The service provides assured capacity, or data rate.
- There is a specified upper bound on the queuing delay through the network. This must be added to the propagation delay, or latency, to arrive at the bound on total delay through the network.
- There are no queuing losses. That is, no packets are lost because of buffer overflow; packets may be lost because of failures in the network or changes in routing paths.

With this service, an application provides a characterization of its expected traffic profile, and the service determines the end-to-end delay that it can guarantee. The guaranteed service is the most demanding service provided by ISA. Because the delay bound is firm, the delay has to be set at a large value to cover rare cases of long queuing delays.

Controlled Load

The key elements of the controlled load service are as follows:

- The service tightly approximates the behavior visible to applications receiving best effort service under unloaded conditions.
- There is no specified upper bound on the queuing delay through the network. However, the service ensures that a very high percentage of the packets do not experience delays that greatly exceed the minimum transit delay.
- A very high percentage of transmitted packets will be successfully delivered.

Best Effort

The risk in an internet that provides QoS for real-time applications is that best effort traffic is crowded out. This is because best effort types of applications are assigned a low priority and their traffic is throttled in the face of congestion and delays. The controlled load service guarantees that the network will set aside sufficient resources so that an application that receives this service will see a network that responds as if these real-time applications were not present and competing for resources.

1.4.4 Queuing Discipline

An important component of an ISA implementation is the queuing discipline used at the routers. The simplest approach that can be used by a router is a first-in, first-out (FIFO) queuing discipline at each output port. As long as the queue is not empty, the router transmits packets from the queue, taking the oldest remaining packet next.

There are several drawbacks to the FIFO queuing discipline:

- No special treatment is given to packets from flows that are of higher priority or are more delay sensitive. If a number of packets from different flows are ready to be forwarded, they are handled strictly in FIFO order.
- If a number of smaller packets are queued behind a long packet, FIFO queuing results in a larger average delay per packet than if the shorter packets were transmitted before the longer packet. In general, flows of larger packets get better service.
- A selfish TCP connection, which ignores the TCP congestion control rules, can crowd out conforming connections. If congestion occurs and one TCP connection fails to back off, other connections along the same path segment must back off more than they would otherwise have to do.

To overcome the drawbacks of FIFO queuing, a number of more complex routing algorithms have been implemented in routers. These algorithms involve the use of multiple queues at each output port and some method of prioritizing the traffic to provide better service.:

- Priority queuing (PQ)
- Custom queuing (CQ)
- Flow-based weighted fair queuing (WFQ)
- Class-based weighted fair queuing (CBWFQ)

For **priority queuing**, each packet is assigned a priority level, and there is one queue for each priority level. In the Cisco implementation, four levels are used: high, medium, normal, and low. Packets not otherwise classified are assigned to the normal priority. PQ can flexibly prioritize according to network protocol, incoming interface, packet size, source/destination address, or other parameters. The queuing discipline gives absolute preference based on priority.

Custom queuing is designed to allow various applications or organizations to share the network among applications with specific minimum throughput or latency requirements. For CQ, there are multiple queues, with each having a configured byte count. The queues are serviced in round-robin fashion. As each queue is visited, a number of packets are dispatched up to the configured byte count. By providing different byte counts for different queues, traffic on each queue is guaranteed a minimum fraction of the overall capacity. Application or protocol traffic can then be assigned to the desired queue.

The remaining queuing algorithms on the preceding list are based on a mechanism known as fair queuing. With simple fair queuing, each incoming packet is placed in the queue for its flow. The queues are serviced in round-robin fashion, taking one packet from each nonempty queue in turn.

The term *weighted fair queuing* (WFQ) is used in the literature to refer to a class of scheduling algorithms that use multiple queues to support capacity allocation and delay bounds. WFQ may also take into account the amount of service requested by each traffic flow and adjust the queuing discipline accordingly.

Flow-based WFQ, which Cisco simply refers to as WFQ, creates flows based on a number of characteristics in a packet, including source and destination addresses, socket numbers, and session identifiers. The flows are assigned different weights to based on IP precedent bits to provide greater service for certain queues.

Class-based WFQ (CBWFQ) allows a network administrator to create minimum guaranteed bandwidth classes. Instead of providing a queue for each individual flow, a class is defined that consists of one or more flows. Each class can be guaranteed a minimum amount of bandwidth.

1.5 DIFFERENTIATED SERVICES

The differentiated services (DiffServ) architecture is designed to provide a simple, easy-to-implement, low-overhead tool to support a range of network services that are differentiated on the basis of performance. Several key characteristics of DiffServ contribute to its efficiency and ease of deployment:

- IP packets are labeled for differing QoS treatment using the existing IPv4 or IPv6 DSField. Thus, no change is required to IP.
- A service level specification (SLS) is established between the service provider (Internet domain) and the customer prior to the use of DiffServ. This avoids the need to incorporate DiffServ mechanisms in applications. Therefore, existing applications need not be modified to use DiffServ. The SLS is a set of parameters and their values that together define the service offered to a traffic stream by a DiffServ domain.
- A traffic conditioning specification (TCS) is a part of the SLS that specifies traffic classifier rules and any corresponding traffic profiles and metering, marking, discarding/shaping rules which are to apply to the traffic stream.
- DiffServ provides a built-in aggregation mechanism. All traffic with the same DiffServ octet is treated the same by the network service. For example, multiple voice connections are not handled individually but in the aggregate. This provides for good scaling to larger networks and traffic loads.
- DiffServ is implemented in individual routers by queuing and forwarding packets based on the DiffServ octet. Routers deal with each packet individually and do not have to save state information on packet flows.

1.5.1 Services

The DiffServ type of service is provided within a DiffServ domain, which is defined as a contiguous portion of the Internet over which a consistent set of DiffServ policies are administered. Typically, a DiffServ domain would be under the control of one administrative entity. The services provided across a DiffServ domain are defined in an SLA, which is a service contract between a customer and the service provider that specifies the forwarding service that the customer should receive for various classes of packets. A customer may be a user organization or another DiffServ domain. Once the SLA is established, the customer submits packets with the DiffServ octet marked to indicate the packet class. The service provider must ensure that the customer gets at least the agreed QoS for each packet class. To provide that QoS, the service provider must configure the appropriate forwarding policies at each router and must measure the performance being provided for each class on an ongoing basis.

A DiffServ framework document lists the following detailed performance parameters that might be included in an SLA:

- Detailed service performance parameters such as expected throughput, drop probability, and latency.

- Constraints on the ingress and egress points at which the service is provided, indicating the scope of the service.
- Traffic profiles that must be adhered to for the requested service to be provided, such as token bucket parameters.
- Disposition of traffic submitted in excess of the specified profile.
- The framework document also gives some examples of services that might be provided:
- Traffic offered at service level A will be delivered with low latency.
- Traffic offered at service level B will be delivered with low loss.
- 90 percent of in-profile traffic delivered at service level C will experience no more than 50 ms latency.
- 95 percent of in-profile traffic delivered at service level D will be delivered.
- Traffic offered at service level E will be allotted twice the bandwidth of traffic delivered at service level F. Traffic with drop precedence X has a higher probability of delivery than traffic with drop precedence Y.

1.5.2 DiffServ Field

Packets are labeled for service handling by means of the 6-bit DSField in the IPv4 header or the IPv6 header. The value of the DSField, referred to as the DiffServ codepoint (DSCP), is the label used to classify packets for differentiated services.

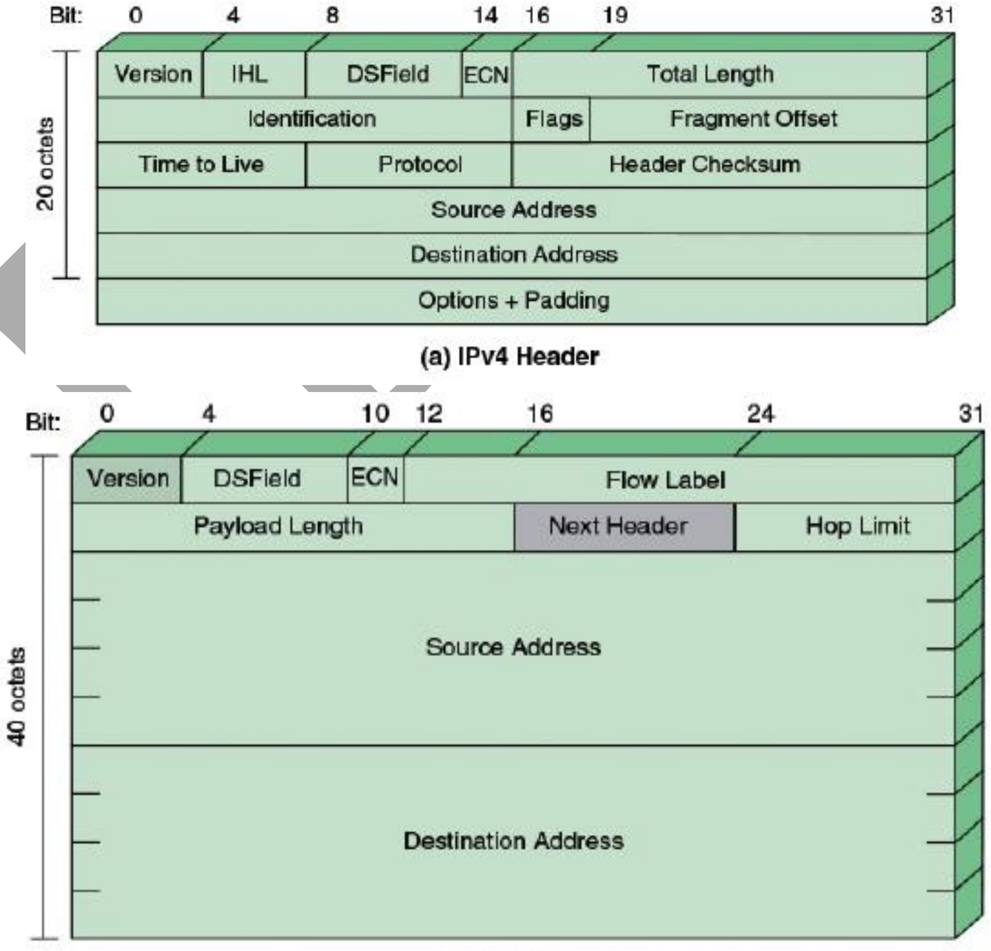

(a) IPv4 Header

(b) IPv6 Header

Figure 1.3

With a 6-bit codepoint, there are in principle 64 different classes of traffic that could be defined. These 64 codepoints are allocated across three pools of codepoints, as follows:

- Codepoints of the form xxxxx0, where x is either 0 or 1, are reserved for assignment as standards.
- Codepoints of the form xxxx11 are reserved for experimental or local use.
- Codepoints of the form xxxx01 are also reserved for experimental or local use.

1.5.3 DiffServ Configuration

Figure 1.4 illustrates the type of configuration envisioned in the DiffServ documents. A DiffServ domain consists of a set of contiguous routers; that is, it is possible to get from any router in the domain to any other router in the domain by a path that does not include routers outside the domain. Within a domain, the interpretation of DS codepoints is uniform, so that a uniform, consistent service is provided.

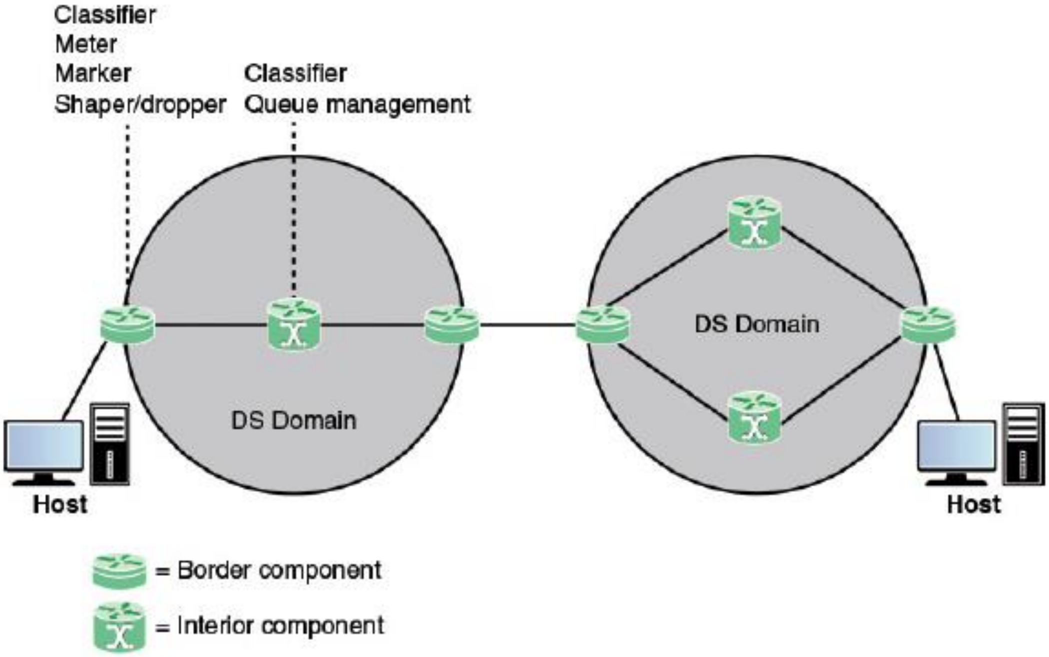

Figure 1.4 DS Domains

Routers in a DiffServ domain are either boundary nodes or interior nodes. Typically, the interior nodes implement simple mechanisms for handling packets based on their DS codepoint values. The DiffServ specifications refer to the forwarding treatment provided at a router as per-hop behavior (PHB). This PHB must be available at all routers, and typically PHB is the only part of DiffServ implemented in interior routers.

The traffic conditioning function consists of five elements:

1. **Classifier**: Separates submitted packets into different classes. This is the foundation of providing differentiated services. A classifier may separate traffic only on the basis of the DS codepoint (behavior aggregate classifier) or based on multiple fields within the packet header or even the packet payload (multifield classifier).
2. **Meter**: Measures submitted traffic for conformance to a profile. The meter determines whether a given packet stream class is within or exceeds the service level guaranteed for that class.

3. **Marker**: Re-marks packets with a different codepoint as needed. This may be done for packets that exceed the profile; for example, if a given throughput is guaranteed for a particular service class, any packets in that class that exceed the throughput in some defined time interval may be re-marked for best effort handling. Also, re-marking may be required at the boundary between two DiffServ domains.
4. **Shaper**: Delays packets as necessary so that the packet stream in a given class does not exceed the traffic rate specified in the profile for that class.
5. **Dropper**: Drops packets when the rate of packets of a given class exceeds that specified in the profile for that class.

1.5.4 DiffServ Operation

Figure 1.5 illustrates the relationship between the elements of traffic conditioning. After a flow is classified, its resource consumption must be measured. The metering function measures the volume of packets over a particular time interval to determine a flow's compliance with the traffic agreement. If the host is bursty, a simple data rate or packet rate may not be sufficient to capture the desired traffic characteristics. A **token bucket** scheme is an example of a way to define a traffic profile to take into account both packet rate and burstiness.

If a traffic flow exceeds some profile, several approaches can be taken. Individual packets in excess of the profile may be re-marked for lower-quality handling and allowed to pass into the DiffServ domain. A traffic shaper may absorb a burst of packets in a buffer and pace the packets over a longer period. A dropper may drop packets if the buffer used for pacing becomes saturated.

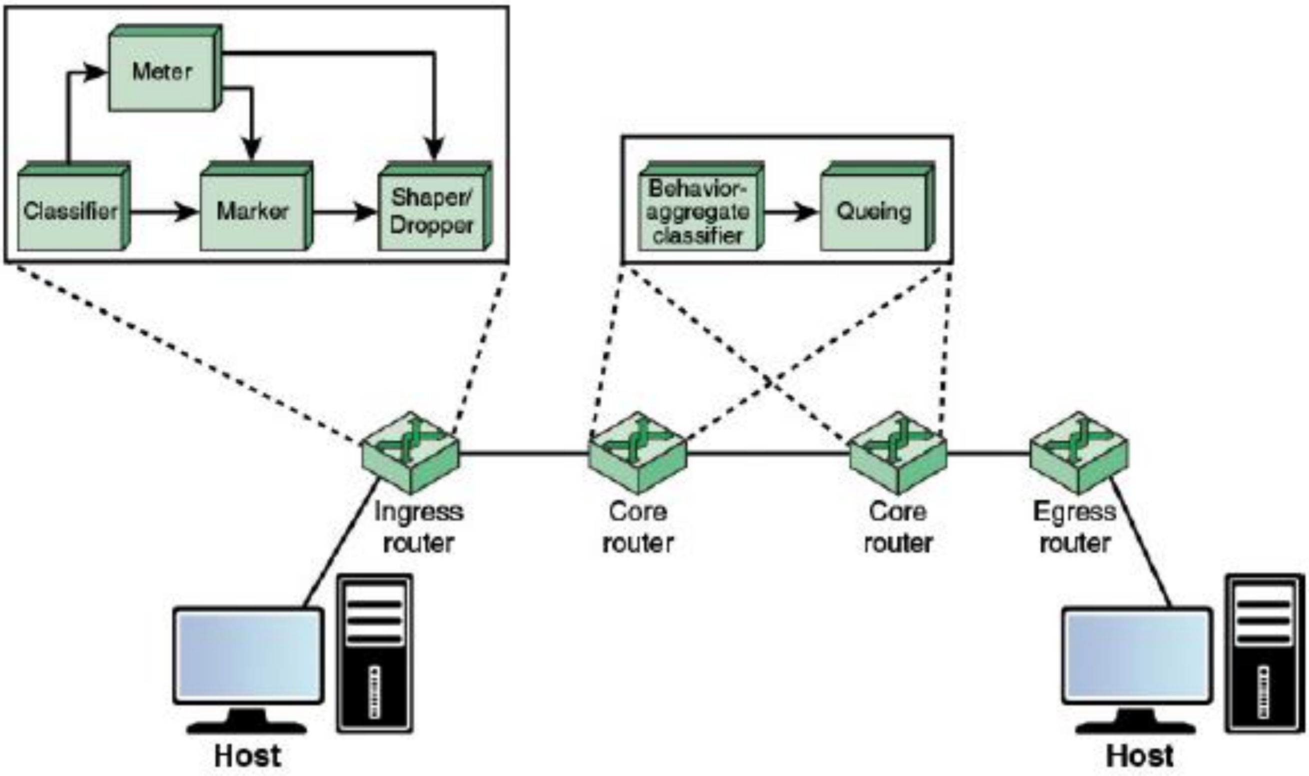

Figure 1.5 DS Functions

1.5.5 Per-Hop Behavior

DiffServ is a general architecture that can be used to implement a variety of services. As part of the DS standardization effort, specific types of PHB need to be defined, which can be associated with specific differentiated services. The four behavior classes are as follows:

1. Default forwarding (DF) for elastic traffic
2. Assured forwarding (AF) for general QoS requirements
3. Expedited forwarding (EF) for real-time (inelastic) traffic
4. Class selector for historical codepoint definitions and PHB requirements

Figure 1.6 shows the DSCP encodings corresponding to the four classes.

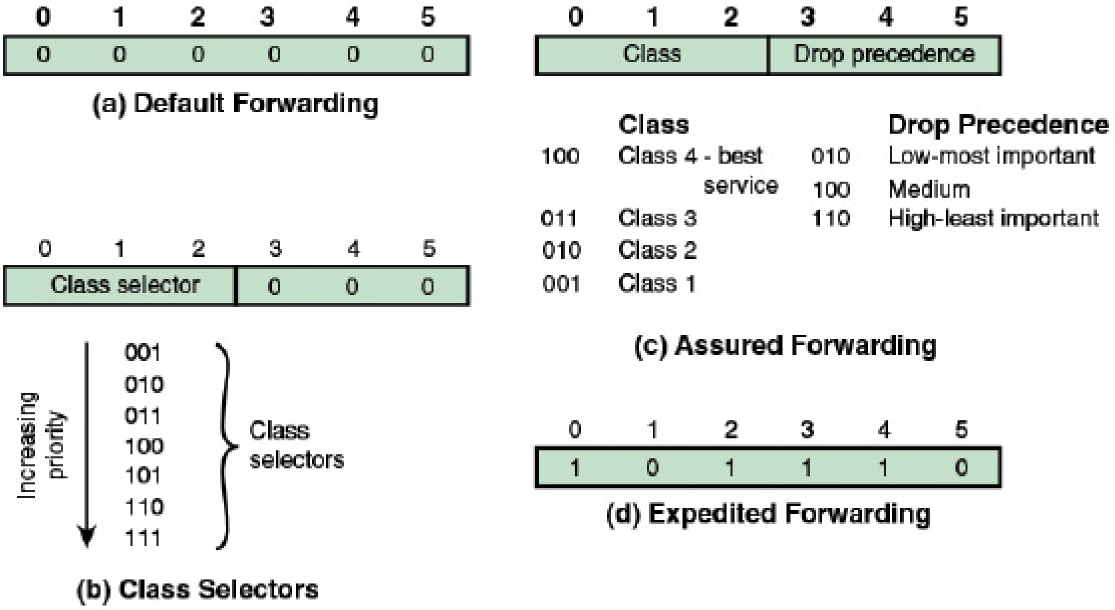

(a) Default Forwarding

(b) Class Selectors

(c) Assured Forwarding

(d) Expedited Forwarding

Figure 1.6 DiffServ Forwarding Behavior Classes and Corresponding DSField Encoding

Default Forwarding PHB

The default class, referred to as default forwarding (DF), is the best effort forwarding behavior in existing routers. Such packets are forwarded in the order that they are received as soon as link capacity becomes available. If other higher-priority packets in other DiffServ classes are available for transmission, the latter are given preference over best effort default packets. Application traffic in the Internet that uses default forwarding is expected to be elastic in nature.

Expedited Forwarding PHB

RFC 3246 defines the expedited forwarding (EF) PHB as a building block for low-loss, low-delay, and low- jitter end-to-end services through DiffServ domains. Therefore, unless the internet is grossly oversized to eliminate all queuing effects, care must be taken in handling traffic for EF PHB to ensure that queuing effects do not result in loss, delay, or jitter above a given threshold.

The EF PHB is designed to configure nodes so that the traffic aggregate has a well-defined minimum departure rate. The general concept outlined in RFC 3246 is this: The border nodes control the traffic aggregate to limit its characteristics (rate, burstiness) to some predefined level. Interior nodes must treat the incoming traffic in such a way that queuing effects do not appear. In general terms, the requirement on interior nodes is that the aggregate's maximum arrival rate must be less than the aggregate's minimum departure rate.

Assured Forwarding PHB

The assured forwarding (AF) PHB is designed to provide a service superior to best effort but one that does not require the reservation of resources within an Internet and does not require the use of detailed discrimination among flows from different users. The AF PHB is more complex than explicit allocation, but it is useful to first highlight the key elements of the explicit allocation scheme:

- Users are offered the choice of a number of classes of service for their traffic. Each class describes a different traffic profile in terms of an aggregate data rate and burstiness.
- Traffic from a user within a given class is monitored at a boundary node. Each packet in a traffic flow is marked out or in based on whether it does or does not exceed the traffic profile.
- Inside the network, there is no separation of traffic from different users or even traffic from different classes. Instead, all traffic is treated as a single pool of packets, with the only distinction being whether each packet has been marked in or out.
- When congestion occurs, the interior nodes implement a dropping scheme in which out packets are dropped before in packets.
- Different users will see different levels of service because they will have different quantities of in packets in the service queues.

1.6 SERVICE LEVEL AGREEMENTS

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

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

Figure 1.1 shows a typical configuration that lends itself to an SLA. In this case, a network service provider maintains an IP-based network. A customer has a number of private networks (for example, LANs) at various sites. Customer networks are connected to the provider via access routers at the access points.

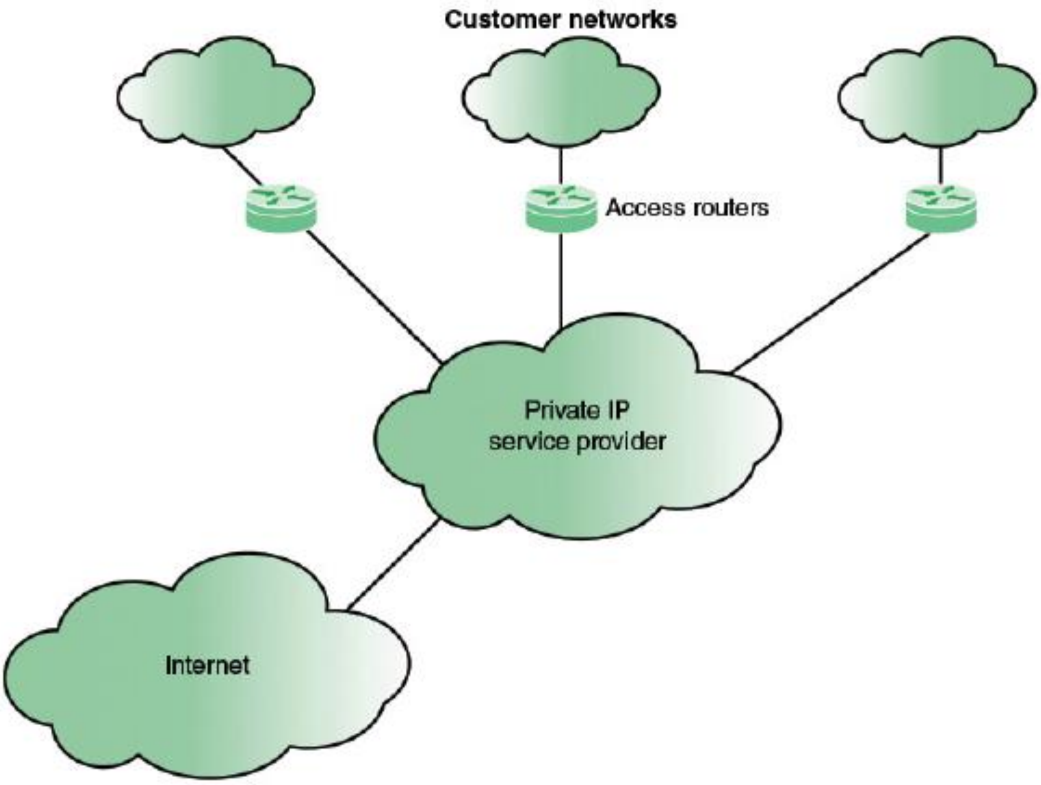

Figure 1.1 Typical Framework for Service Level Agreement

1.1 IP PERFORMANCE METRICS

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

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

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

- Supporting capacity planning and troubleshooting of large complex internets.
- Encouraging competition by providing uniform comparison metrics across service providers.

- Supporting Internet research in such areas as protocol design, congestion control, and QoS.
- Verification of SLAs.

These metrics are defined in three stages:

1. **Singleton metric**: The most elementary, or atomic, quantity that can be measured for a given performance metric. For example, for a delay metric, a singleton metric is the delay experienced by a single packet.
2. **Sample metric**: A collection of singleton measurements taken during a given time period. For example, for a delay metric, a sample metric is the set of delay values for all the measurements taken during a one-hour period.
3. **Statistical metric**: A value derived from a given sample metric by computing some statistic of the values defined by the singleton metric on the sample. For example, the mean of all the one-way delay values on a sample might be defined as a statistical metric.

The measurement technique can be either active or passive.

- **Active techniques** require injecting packets into the network for the sole purpose of measurement. There are several drawbacks to this approach. The load on the network is increased. This, in turn, can affect the desired result.
- **Passive techniques** observe and extract metrics from existing traffic. This approach can expose the contents of Internet traffic to unintended recipients, creating security and privacy concerns. So far, the metrics defined by the IPPM working group are all active.

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

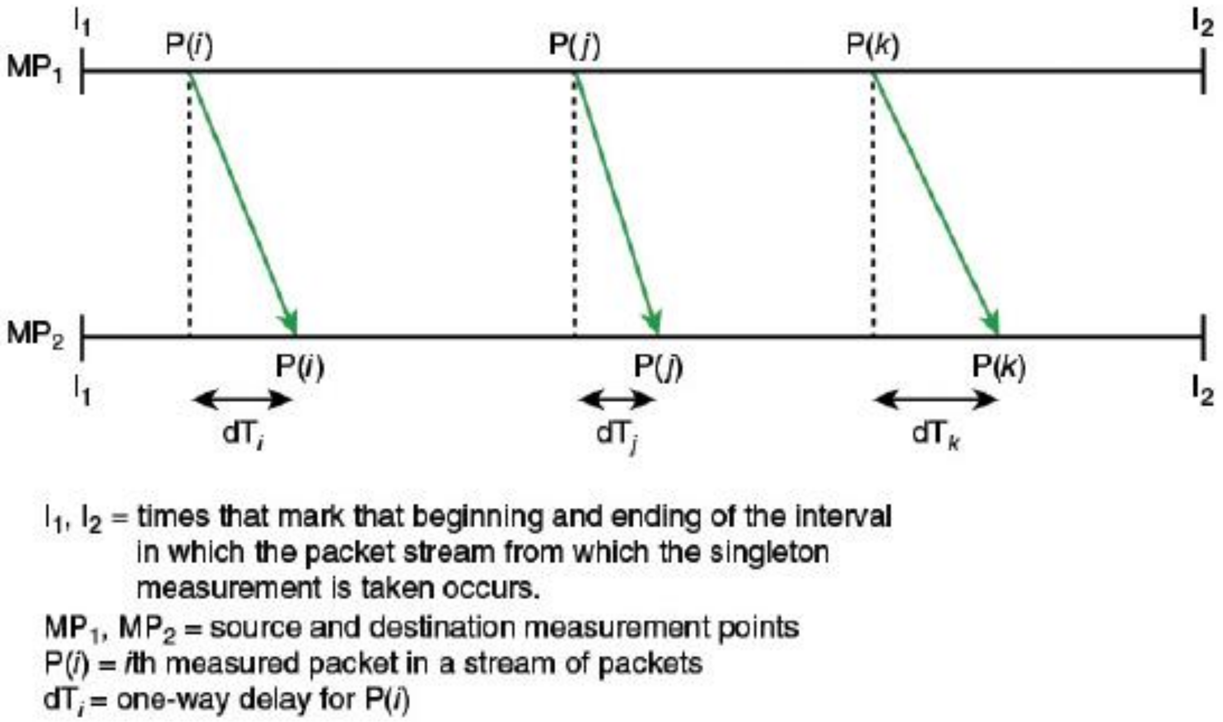
I1
I2
MP1
P(i)
P(j)
P(k)
MP2
I1
P(i)
P(j)
P(k)
I2
dTi
dTj
dTk
I1, I2 = times that mark that beginning and ending of the interval in which the packet stream from which the singleton measurement is taken occurs.
MP1, MP2 = source and destination measurement points
P(i) = ith measured packet in a stream of packets
dTi = one-way delay for P(i)

Figure 1.8 Model for Defining Packet Delay Variation

Figure 1.8 illustrates the packet delay variation metric. This metric is used to measure jitter, or variability, in the delay of packets traversing the network. The singleton metric is defined by selecting two packet measurements and measuring the difference in the two delays. The statistical measures make use of the absolute values of the delays.

1.8 OPENFLOW QOS SUPPORT

OpenFlow offers two tools for implementing QoS in data plane switches.

1.8.1 Queue Structures

An OpenFlow switch provides limited QoS support through a simple queuing mechanism. One or more queues can be associated with a port. Queues support the ability to provide minimum data rate guarantees and maximum data rate limits. Queue configuration takes place outside the OpenFlow protocol, either through a command-line tool or through an external dedicated configuration protocol.

A data structure defines each queue. The data structure includes a unique identifier, port this queue is attached to, minimum data rate guaranteed, and maximum data rate. Counters associated with each queue capture the number of transmitted bytes and packets, number of packets dropped because of overrun, and the elapsed time the queue has been installed in the switch.

The OpenFlow Set-Queue action is used to map a flow entry to an already configured port. Thus, when an arriving packet matches a flow table entry, the packet is directed to a given queue on a given port.

1.8.2 Meters

A meter is a switch element that can measure and control the rate of packets or bytes. Associated with each meter is a set of one or more bands. If the packet or byte rate exceeds a predefined threshold, the meter triggers the band. The band may drop the packet, in which case it is called a **rate limiter**. Other QoS and policing mechanisms can be designed using meter bands. Each meter is defined by an entry in the meter table for a switch. Each meter has a unique identifier. Meters are not attached to a queue or a port; rather, a meter can be invoked by an instruction from a flow table entry. Multiple flow entries can point to the same meter.

Figure 1.9 shows the structure of a meter table entry and how it is related to a flow table entry.

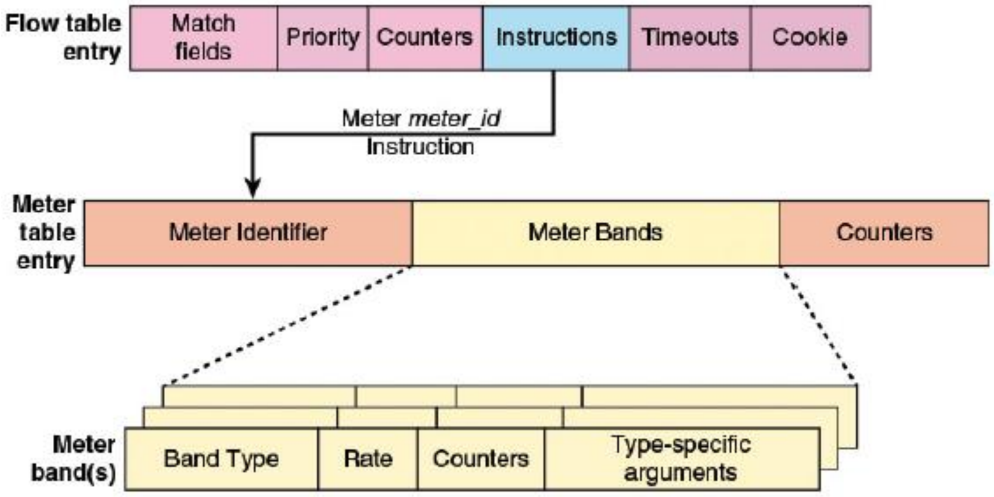

Figure 1.9 OpenFlow QoS-Related Formats

A flow table entry may include a meter instruction with ameter_id as an argument. Any packet that matches that flow entry is directed to the corresponding meter. Within the meter table, each entry consists of three main fields:

- **Meter identifier:** A 32-bit unsigned integer uniquely identifying the meter.
- **Meter bands:** An unordered list of one or more meter bands, where each meter band specifies the rate of the band and the way to process the packet.
- **Counters:** Updated when packets are processed by a meter. These are aggregate counters. That is, the counters count the total traffic of all flows, and do not break the traffic down by flow.

Each band has the following structure:

- **Band type:** drop **or** dscp remark**.**
- **Rate:** Used by the meter to select the meter band, defines the lowest rate at which the band can apply.
- **Counters:** Updated when packets are processed by a meter band.
- **Type specific arguments:** Some band types may have optional arguments. Currently, the only optional argument is for the ds cp remark band type, specifying the amount of drop in precedence.

The meter triggers a meter band if the packet rate or byte rate passing through the meter exceed a predefined threshold. A band of type drop drops packets when the band's rate is exceeded. This can be used to define a rate limiter band.

1.9 QoE: User Quality of Experience

WHY QOE?

Before the advent of the public Internet, video content delivery was a monopoly of content publishers who delivered their products and services over closed video delivery systems built and managed by cable and satellite TV operators. The operators owned and operated the entire distribution chain as well as the video reception devices (set-top boxes) in the home. These closed networks and devices were under the full control of these operators and were designed, deployed, provisioned, and optimized specifically to deliver high-quality video to consumers.

Figure 1.10 shows an abstraction of the typical satellite TV end-to-end delivery chain. In practice, however, such content delivery and distribution chains are made up of very complex integrations of applications and systems.

As the illustration shows, the traffic scheduling system provides audio and video (A/V) content via the play-out system to be encoded and aggregated into a single MPEG transport stream (TS). Together with the program specific information (PSI), the transport stream is transmitted to the subscriber's set-top box (STB) via a satellite.

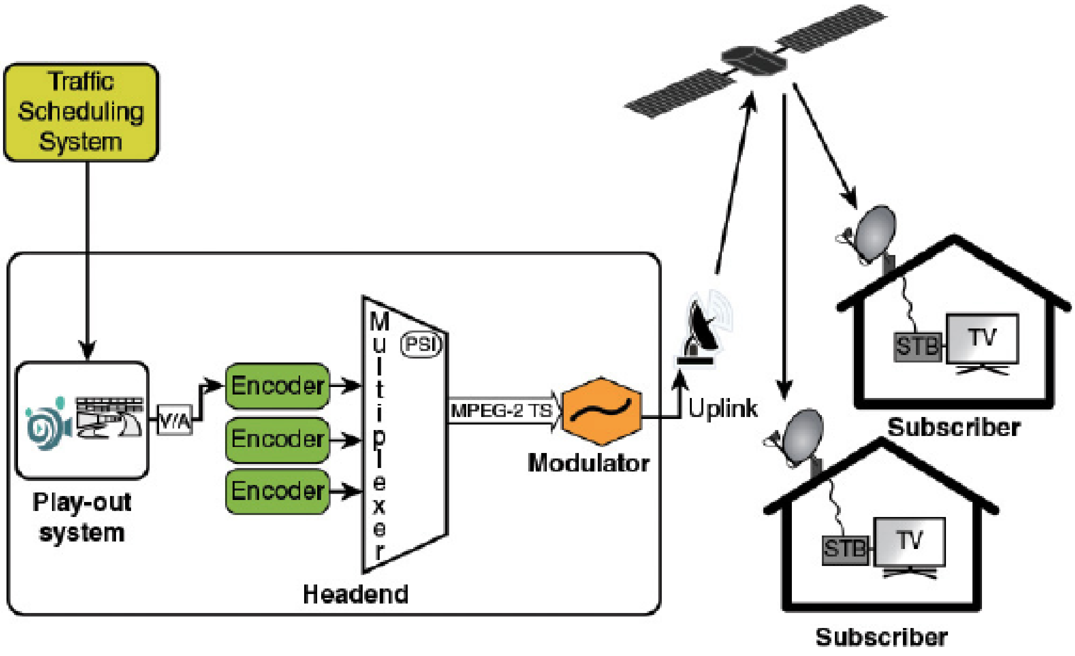

Figure 1.10 An Abstraction of a Content Distribution Network Using a Typical Satellite TV Distribution Network

Online Video Content Delivery

Video delivery over the Internet takes a different approach. Because numerous subnetworks and devices that constitute the Internet are situated in varied geographical locations, video streams reach the user by traversing through uncharted territories, as illustrated in Figure 1.11. With this arrangement, the guaranteeing of a good network performance is often a very challenging task.

Internet service providers (ISPs) do not own the entire content distribution network, and the risk of quality degradations is high. The access network may consist of coax, copper, fiber, or wireless (fixed and mobile) technology. Issues such as packet delay, jitter, and loss may plague such networks.

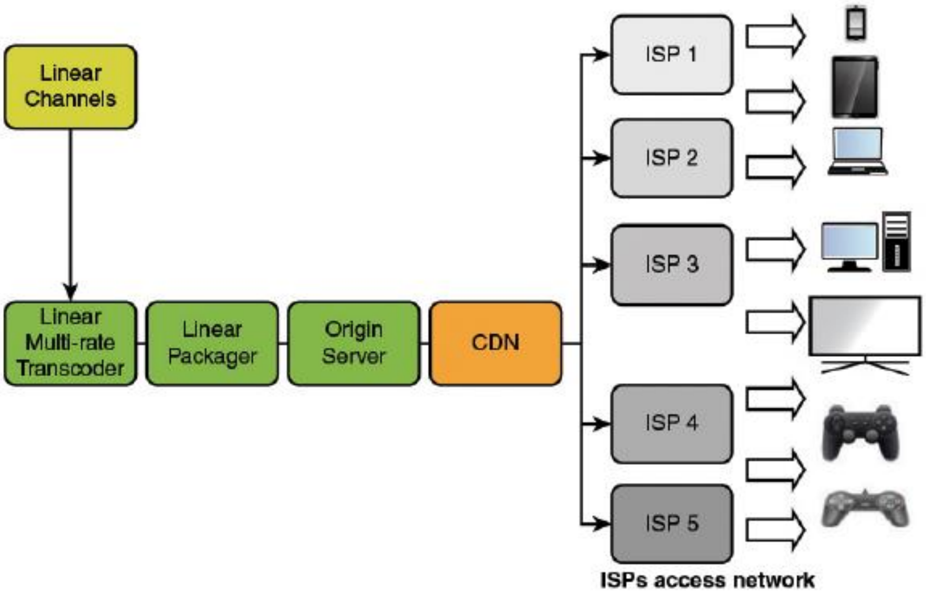

Figure 1.11 An Abstraction of a Content Distribution Network Using the Public Internet Distribution Network

The growth and expansion of the Internet over the past couple of decades has led to an equally huge growth in the availability of network-enabled video streaming services. Giant technological strides have also been made in the development of network access devices.

With the current popularity of these services, providers need to ensure that user experiences are comparable to what the users would consider to be their reference standards. Users' standards are often influenced by the typically high video quality experience with the older technology, that is, those offered by the cable and satellite TV operators. User expectations can also be influenced by capabilities that currently can only be adequately offered by broadcast TV. These capabilities include the following:

Trick mode functionalities, which are features of video streaming systems that mimic visual feedback given during fast-forward and rewind operations.

Contextual experiences across multiple screens, which includes the ability to pause viewing on one screen and switch to another, thus letting users take the video experience with them on the go.

The proliferation of different types of access devices further highlights the importance of QoE frameworks. As an illustration, the QoE for a user watching a news clip on a PDA will most likely differ from another user watching that same news clip on a 3G mobile phone. This is because the two terminals come with different display screens, bandwidth capabilities, frame rates, codecs, and processing power. Therefore, delivering multimedia content or services to these two terminal types, without carefully thinking about the users' quality expectations or requirements for these terminal types, might lead to service overprovisioning and network resource wastage.

Informally, QoE refers to the user perception of a particular service. QoE needs to be one of the central metrics employed during the design and management of networks, content delivery systems, and other engineering processes.

1.10 SERVICE FAILURES DUE TO INADEQUATE QOE CONSIDERATIONS

The stereoscopic 3D TV service is often cited as a prime example of a service that was a spectacular commercial failure because it had very poor QoE ratings.

In 2010, broadcasters such as Disney, Foxtel, BBC, and Sky began actively making 3D content delivery available as a service to their customers as a premium service experience. Indeed, each of these broadcasters rolled out their own dedicated 3D television channels. Within five years, all of them except Sky had to terminate their operations.

A number of factors contributed to the failure of these services.

- The first was the general unavailability of "wow video content" (that is, content that users are most likely to find exciting or take much interest in).
- The second was the need to wear special 3D glasses even when using these services in a home environment.
- Third, because broadcasters were initially in a rush to deploy the 3D TV technology, content was produced by inexperienced creators using inadequate systems and tools. This

resulted in a great deal of poorly produced 3D content, which may have alienated the early subscribers.

1.11 QOE-RELATED STANDARDIZATION PROJECTS

Because the field of QoE has been growing rapidly, a number of projects have been initiated to address issues relating to best practices and standards. These projects have been aimed at preventing commercial. Table 1.1 summarizes the prominent ones amongst these project initiatives.

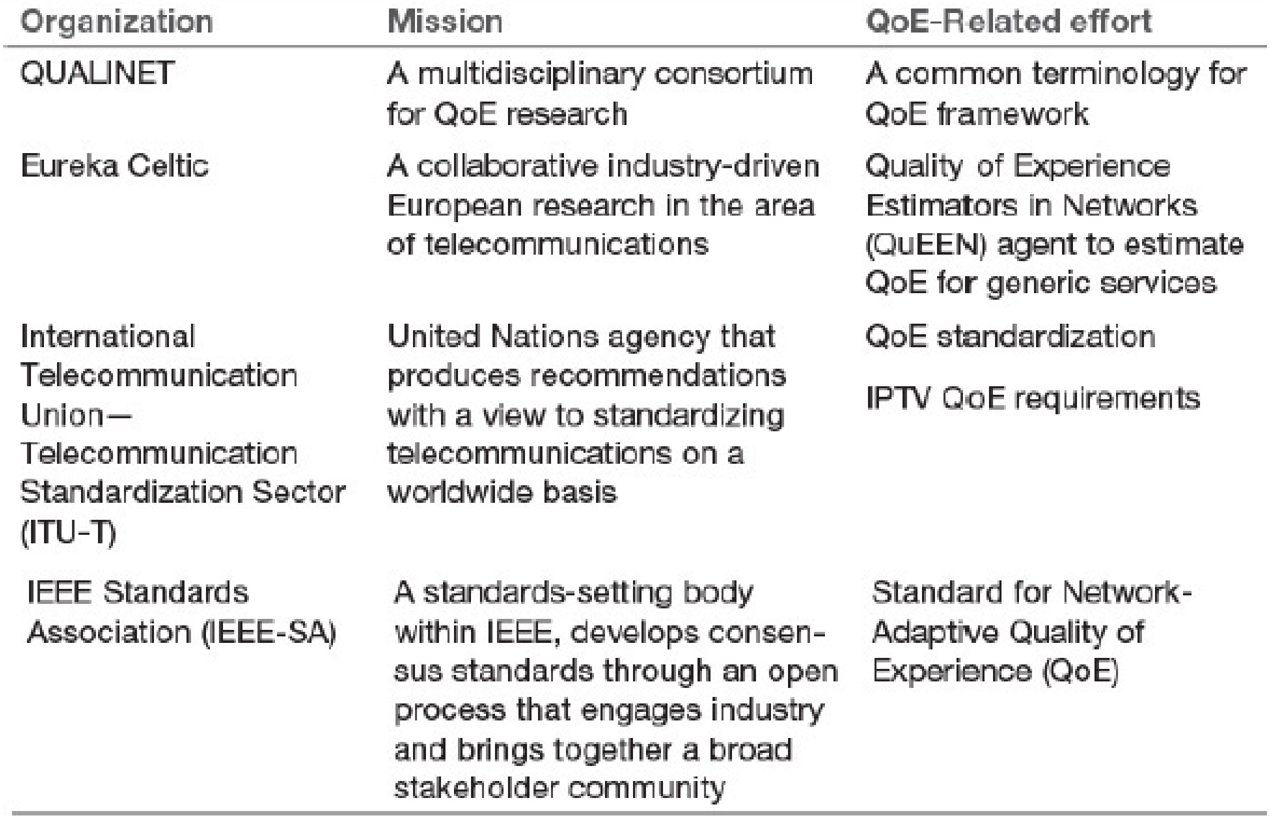

Organization	Mission	QoE-Related effort
QUALINET	A multidisciplinary consortium for QoE research	A common terminology for QoE framework
Eureka Celtic	A collaborative industry-driven European research in the area of telecommunications	Quality of Experience Estimators in Networks (QuEEN) agent to estimate QoE for generic services
International Telecommunication Union—Telecommunication Standardization Sector (ITU-T)	United Nations agency that produces recommendations with a view to standardizing telecommunications on a worldwide basis	QoE standardization IPTV QoE requirements
IEEE Standards Association (IEEE-SA)	A standards-setting body within IEEE, develops consensus standards through an open process that engages industry and brings together a broad stakeholder community	Standard for Network-Adaptive Quality of Experience (QoE)

Table 1.1 QoE Initiatives and Projects

1.12 DEFINITION OF QUALITY OF EXPERIENCE

Definition of Quality

Quality is the resulting verdict produced by a user after he/she has carried a “comparison and judgment” process on an observable occurrence or event.

This process comprises the following key sequential steps:

- Perception of the event
- Reflection on the perception
- Description of the perception
- Evaluation and description of the result or outcome

Thus, quality is evaluated in terms of the degree to which the user’s needs have been fulfilled within the context of the event. The result of this evaluation is usually referred to as the quality score (or rating) if it is presented with reference to a scale.

Definition of Experience

Experience is an individual's description of a stream of perceptions, and his/her interpretation of one or multiple events. An experience might result from an encounter with a system, service, or an artifact.

Quality Formation Process

As shown in Figure 1.12, there are two distinct subprocess paths to the formation of a quality score: the perception path and the reference path.

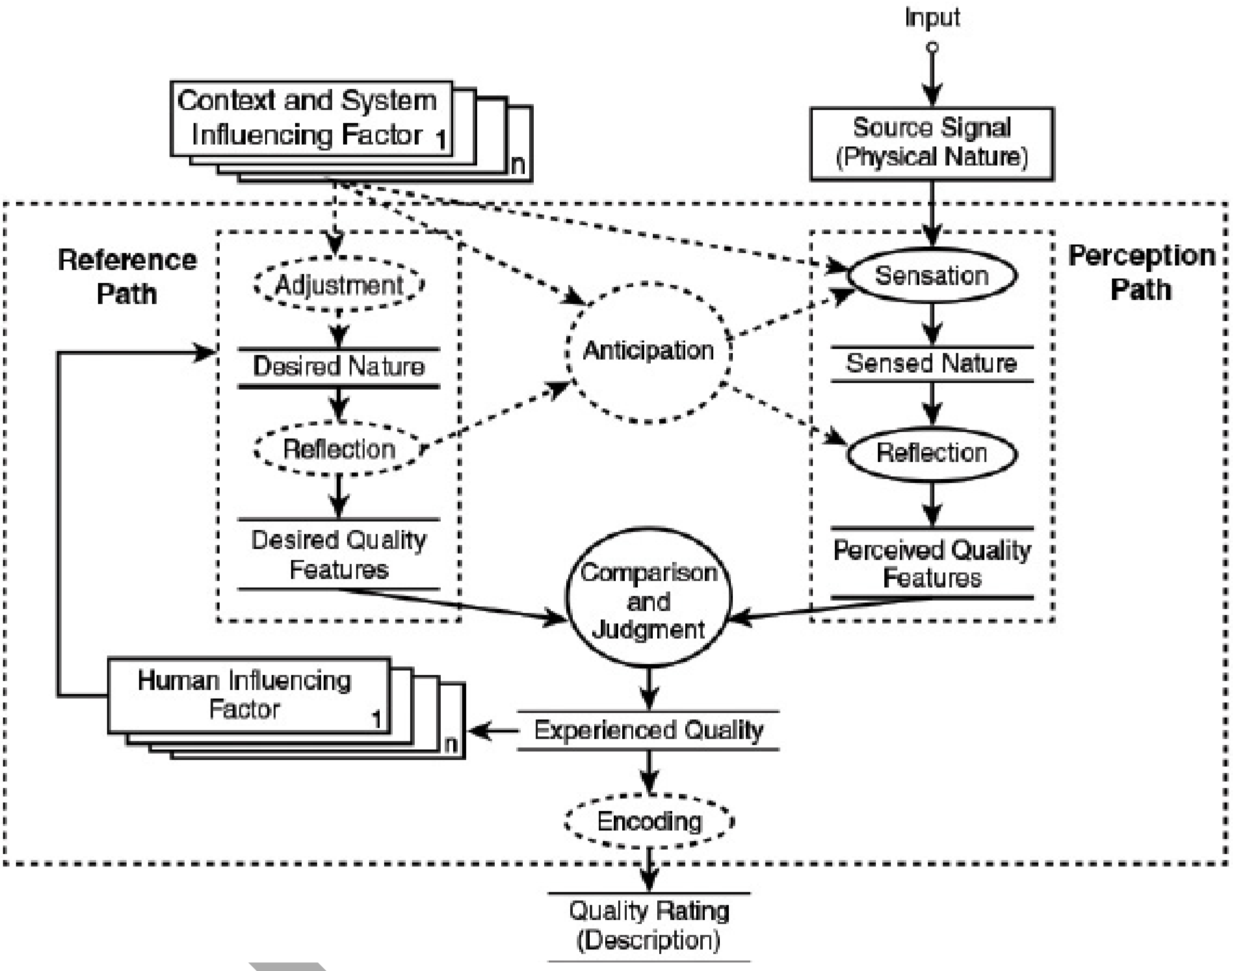

FIGURE 1.12 A Schematic Illustration of the Quality Formation Process from an Individual Point of View.

The reference path reflects the temporal and contextual nature of the quality formation process. This path is influenced by memories of former experienced qualities, as indicated by the arrow from experienced quality to the reference path.

The perception path is characterized by the physical input signal, which is to be assessed, reaching the sensory organs of the observer. This physical event is processed through low-level perceptual processes into a perceived feature within the constraints of the reference path. This perceived feature undergoes a reflection process, which interprets these sensory features through **cognitive** processing.

Definition of Quality of Experience

Combining the concepts and definitions from the preceding sections, the definition of QoE that reflects broad industry and academic consensus is as follows:

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

1.13 QOE STRATEGIES IN PRACTICE

Key findings from QoE-related projects show that for many services, multiple QoS parameters contribute toward the overall user's perception of quality. This has resulted in the emergence of the concept of the QoE/QoS layered approach.

The QoE/QoS Layered Model

The QoE/QoS layered approach does not ignore the QoS aspect of the network, but instead, user and service level perspectives are complementary, as shown in Figure 1.13.

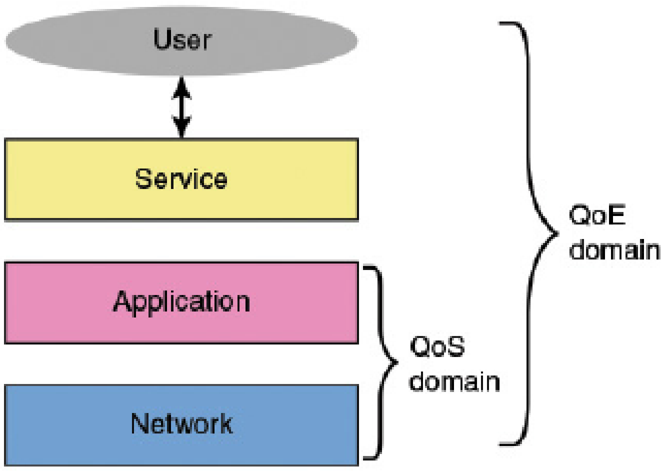

FIGURE 1.13 QoE/QoS Layered Model with the Domains of Interest for the Frameworks

The levels in the layered approach are as follows:

- **User**: The user interacts with the service. It is their degree of delight or annoyance from using the service that is to be measured. Being linked to human perception, QoE is hard to describe in a quantitative way, and it varies from person to person. The complexities of QoE at the user level stem from the differences between individual user characteristics, of which some might be time-varying, whereas others are of a relatively stable nature. The current practice in any QoE measurement is to identify and control for the relatively stable characteristics of a user in a way that is satisfactory to at least a large proportion of the potential user group.
- **Service**: The service level provides a virtual level where the user's experience of the overall performance of the service can be measured. It is the interface where the user interacts with the service (for example, the visual display to the user). It is also where tolerance thresholds are measured. As an illustration, the QoE measures from the user

perspective for streaming applications could be startup time, audio/visual quality, channel change delay, and buffering interruptions.

- **Application-level QoS (AQoS)**: AQoS deals with the control of application-specific parameters such as content resolution, bit rate, frame rate, color depth, codec type, layering strategy, and sampling rate. The network capacity often dictates the bandwidth that will be allocated to a service for transmission. Because of this fixed underlying resource, some parameters at the application level are usually adjusted and controlled to achieve a desired quality level.
- **Network-level QoS (NQoS)**: This level is concerned with the low-level network parameters such as service coverage, bandwidth, delay, throughput, and packet loss. There are a number of ways in which network-level QoS parameters impact QoE. One such way is via network delay, which impacts QoE especially for interactive services. For instance, the interactive nature of web browsing that requires multiple retrieval events within a certain window of time might be affected by delay variations of the network. Voice over IP (VoIP) services might have stringent response-time demands, whereas e-mail services might tolerate much longer delays.

Although the trade-offs between quality and network capacity may begin with application-level QoS because of network capacity considerations, an understanding of the user requirements at the service level (that is, in terms QoE measures) would enable a better choice of application-level QoS parameters to be mapped onto the network-level QoS parameters.

1.14 FACTORS INFLUENCING QOE

QoE must be studied and addressed by taking into account both technical and nontechnical factors. Many factors contribute to producing a good QoE. Here, the key factors are as follows:

- **User demographics**: The context of demographics herein refers to the relatively stable characteristics of a user that might have an indirect influence on perception, and intimately affects other technical factors to determine QoE. The grouping of users was based on demographic characteristics such as their attitudes toward adoption of new technologies, socio-demographic information, socioeconomic status, and prior knowledge. Cultural background is another user demographic factor that might also have an influence on perception because of cultural attitude to quality.
- **Type of device**: Different device types possess different characteristics that may impact on QoE. An application designed to run on more than one device type, for example on a connected TV device such as Roku and on an iOS device such as an iPhone, may not deliver the same QoE on every device.
- **Content:** Content types can range from interactive content specifically curated according to personal interests, to content that is produced for linear TV transmission. Studies have suggested that people tend to watch video on-demand (VoD) content with a higher level of engagement than its competing alternative, linear TV. This may be because users will make an active decision to watch specific VoD content, and as a result, give their full attention to it.
- **Connection type:** The type of connection used to access the service influences users' expectations and their QoEs. Users have been found to have lower expectations when using 3G connections in contrast to a wire line connection even when the two

connection types were identical in terms of their technical conditions. Users have also been found to lower their expectations considerably, and are more tolerant to visual impairments, on small devices.

- **Media (audio-visual) quality:** This is a significant factor affecting QoE, as it is the part of a service that is most noticeable by the user. The overall audio and video quality appears to be content dependent. For less- complex scenes (for example, head and shoulder content), audio quality is slightly more important than video quality. In contrast, for high-motion content, video quality tends to be significantly more important than audio quality.
- **Network:** Content delivery via the Internet is highly susceptible to the effects of delays, jitter, packet loss, and available bandwidth. Delay variation results in the user experiencing frame freeze and the lack of lip synchronization between what is heard (audio) and what is seen (video). Although video content can be delivered using a number of Internet protocols, not all of them are reliable. However, content delivery is guaranteed using TCP/IP. Nevertheless, bad network conditions degrade QoE because of increased rebuffering and increased interruptions in playback. Rebuffering interruptions in IP video playback is seen to be the worst degradation on user QoE and should be avoided at the cost of startup delay.
- **Usability:** Another QoE factor is the amount of effort that is required to use the service. The service design must render good quality without a great deal of technical input from the user.
- **Cost:** The long-established practice of judging quality by price implies that expectations are price dependent. If the tariff for a certain service quality is high, users may be highly sensitive to any quality degradations.

1.15 MEASUREMENTS OF QOE

QoE measurement techniques evolved through the adaptation and application of psychophysics methods during the early stages of television systems. We will see three QoE measurement methods:

1. Subjective Assessment
2. Objective Assessment
3. End-User Device Analytics

1.15.1 Subjective Assessment

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

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

- **Design and define test matrix:** Once the service has been characterized, the QoS factors that affect the QoE measures can be identified. For instance, the video quality in streaming services might be directly affected by network parameters such as bandwidth, packet loss, and encoding parameters such as frame rate, resolution, and codec. The capability of the rendering device will also play a significant role in terms of screen size and processing power. However, testing such a large combination of parameters may not be feasible.
- **Specify test equipment and materials:** Subjective tests should be designed to specify test equipment that will allow the test matrix to be enforced in a controlled fashion. For instance, to assess the correlation between NQoS parameters and the perceived QoE in a streaming application, at least a client device and a streaming server separated by an emulated network are needed.
- **Identify sample population:** A representative sample population is identified, possibly covering different classes of users categorized by the user demographics that are of interest to the experimenter. Depending on the target environment for the subjective test, at least 24 test subjects has been suggested as the ideal number for a controlled environment (for example, a laboratory) and at least 35 test subjects for a public environment. Fewer subjects may be used for pilot studies to indicate trending. The use of crowdsourcing in the context of subjective assessment is still nascent, but it has the potential to further increase the size of the sample population and could reduce the completion time of the subjective test.
- **Subjective methods**: Several subjective assessment methodologies exist within the industry recommendations. However, in most of them, the typical recommendation is for each test subject to be presented with the test conditions under scrutiny along with a set of rating scales that allows the correlation of the users' responses with the actual QoS test conditions being tested. There are several rating scales, depending on the design of the experiment.
- **Analysis of results**: When the test subjects have rated all QoS test conditions, a post-screening process might be applied to the data to remove any erroneous data from a test subject that appears to have voted randomly. Depending on the design of the experiment, a variety of statistical approaches could be used to analyze results. The simplest and the most common quantification method is the mean opinion score (MOS), which is the average of the opinions collected for a particular QoS test condition. The results from subjective assessment experiments are used to quantify QoE, and to model the impacts of QoS factors. However, they are time-consuming, expensive to carry out, and are not feasible for real-time in-service monitoring.

1.15.2 Objective Assessment

For objective assessment of QoE, computational algorithms provide estimates of audio, video, and audiovisual quality as perceived by the user. Each objective model targets a specific service type. The goal of any objective model is to find the optimum fit that strongly correlates with data obtained from subjective experiments. A methodology to obtain objective QoE data might consist of the following phases:

- **Database of subjective data**: A starting point might be the collection of a group of subjective datasets as this could serve as benchmark for training and verifying the performance of the objective model. A typical example of one of these datasets might be the subjective QoE data generated from well-established subjective testing procedures.
- **Preparation of objective data**: The data preparation for the objective model might typically include a combination of the same QoS test conditions as found in the subjective datasets, as well as other complex QoS conditions. A variety of preprocessing procedures might be applied to the video data prior to training, and refinement of the algorithm.
- **Objective methods**: There are various algorithms in existence that can provide estimates of audio, video, and audiovisual quality as perceived by the user. Some algorithms are specific to a perceived quality artifact, while others can provide estimates for a wider scope of quality artifacts. Examples of the perceived artifacts might include blurring, blockiness, unnatural motion, pausing, skipping, rebuffering, and imperfect error concealment after transmission errors.
- **Verification of results**: After the objective algorithm has processed all QoS test conditions, the predicted values might benefit from a post-screening process to remove any outliers; this is the same concept applied to the subjective datasets. The predicted values from the objective algorithm might be in a different scale as compared to the subjective QoE datasets.
- **Validation of objective model**: The objective data analysis might be evaluated with respect to its prediction accuracy, consistency, and linearity by using a different subjective dataset. It is worth noting that the performance of the model might depend on the training datasets and the verification procedures. The Video Quality Experts Group (VQEG) validates the performance of objective perceptual models.

1.15.3 End-User Device Analytics

End-user device analytics is yet another alternative method of QoE measurement. Real-time data such as the connection time, bytes sent, and average playback rate are collected by the video player application for each video viewing session and fed back to a server module where the data is pre-aggregated and then turned into actionable QoE measures. Some of the metrics reported for per-user and aggregate viewing sessions include startup delay, rebuffering delays, average bit rates, and the frequency of bit rate switches.

Operators may be inclined to associate viewer engagement levels with their QoE because good QoEs usually make viewers less likely to abandon a viewing session. The definition of viewer engagement may have different meanings for different operators and context. First of all, operators might like to know which viewer engagement metrics affect QoE the most to guide the design of the delivery infrastructures. Second, they might also like to quickly identify and resolve service outages, and other quality issues. A minute of encoder glitch could replicate throughout the ISPs, and the various delivery infrastructures, and affect all their customers. Operators might like to know the scale of this impact, and how it affects users' engagement.

Finally, they would like to understand their customers' demographics (connection methods, type of device, bit rates of the consumed asset) within a demographic region so that resources can be strategically dimensioned.

1.16 APPLICATIONS OF QOE

The practical applications of QoE can be grouped into two areas based on the main usage.

1. **Service QoE monitoring:** Service monitoring allows the support teams (for example, service provider and network operator) to continually monitor the quality experienced by the end users of the service. A service alert message might be sent to the support teams when QoE falls below a certain threshold value, as this will allow the support teams to quickly identify and resolve service outages and other QoE issues.
2. **QoE-centric network management:** The ability to control and optimize the user experience when QoE degradation issues arise is the holy grail of QoE network management. Given the multidimensional aspect of the overall QoE (such as the network-level conditions of the subnetworks, application-level QoS, device capability, and user demographics), a typical challenge lies in providing actionable QoE information feedback to the network or service provider.

Two approaches in which QoE-centric network management can be exploited are as follows:

- In the first approach, a set of QoS measurement values together with the appropriate assumptions, are used in computing the expected QoE for a user.
- In the second approach, which is somewhat the opposite of the first, a target QoE for a user together with the appropriate assumptions is used to produce estimates of the required QoS values.

The first approach can be taken by a service provider, who can provide a range of QoS offerings with an outline of the QoE that the customer might reasonably expect.

The second approach can be taken by a customer who defines the required QoE, and then determines what level of service will meet that need.

Figure 1.14 illustrates a scenario where the user can make a selection from a range of services, including the required level of service (SLA). By contrast to the purely QoS-based management, the SLA here is not expressed in terms of raw network parameters. Instead, the user indicates a QoE target; it is the service provider that maps this QoE target together with the type of service selected, onto QoS demands.

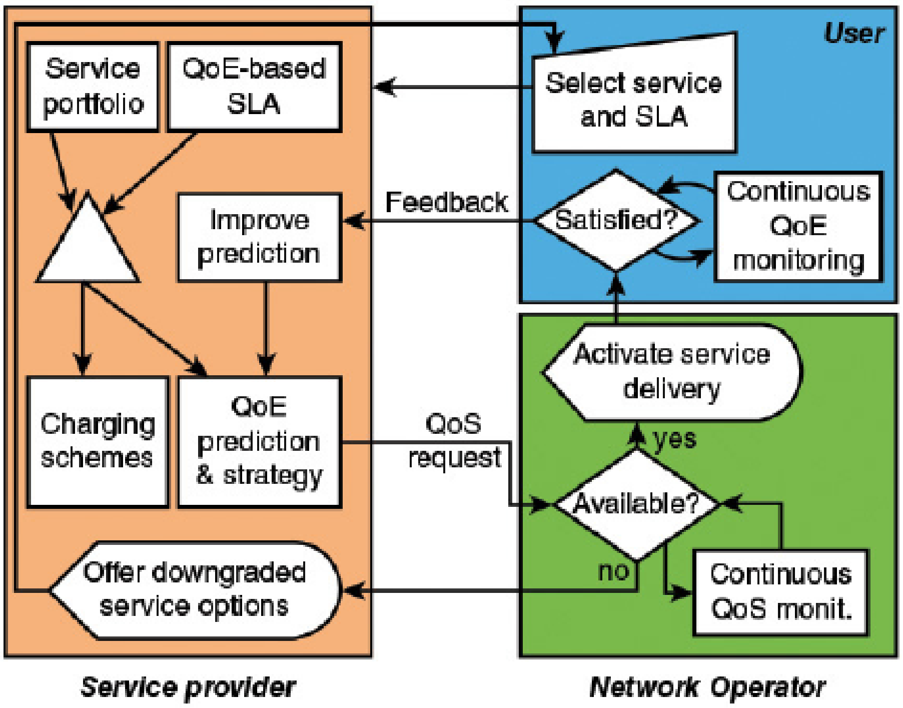

FIGURE 11.5 QoE-Centric Network Management

Figure 1.14 QoE-Centric Network Management

The service provider selects the appropriate quality prediction model and management strategy (for example, minimize network resource consumption) and forwards a QoS request to the operator. It is possible that the network cannot sustain the required level of QoS, making it impossible to deliver the requested QoE. This situation leads to a signal back to the user, prompting a reduced set of services/QoE values.

Unit 4: Chapter 2

Network Design Implications of QoS and QoE

Unit Structure

8.0 OBJECTIVES:

- Translate metrics from QoS to QoE domain.
- Select the appropriate QoE/QoS mapping model for a given operational situation.
- Deploy QoE-centric monitoring solutions over a given infrastructure.
- Deploy QoE-aware applications over QoE-centric infrastructure.

8.1 INTRODUCTION:

QoE/QoS mapping model is a function that transforms metrics from QoS to QoE domains.

8.2 CLASSIFICATION OF QOE/QOS MAPPING MODELS

QoE/QoS mapping models can be classified according to their inputs into three categories:

1. Black-box media-based models
2. Glass-box parameter-based models
3. Gray-box parameter-based models

8.2.1 Black-Box Media-Based QoS/QoE Mapping Models

Black-box media-based quality models rely on the analysis of media gathered at system entrance and exit. Hence, they account implicitly for the characteristics of examined media processing system. They are classified into two categories:

a-Double-sided or full-reference quality models: They use as inputs the clean stimulus and the corresponding degraded stimulus. They compare the clean and degraded stimulus in a *perceptual domain* that accounts for psychophysics capability of human sensory system. The perceptual domain is a transformation of traditional physical temporal and frequency domains performed according to characteristics of users perceptions. Basically, the larger the perceptual distance, the greater the degradation level. This model needs to align clean and degraded stimulus because the comparison is made on per-block basis. The stimulus alignment should be realized autonomously, that is, without adding extra control information describing stimulus structure.

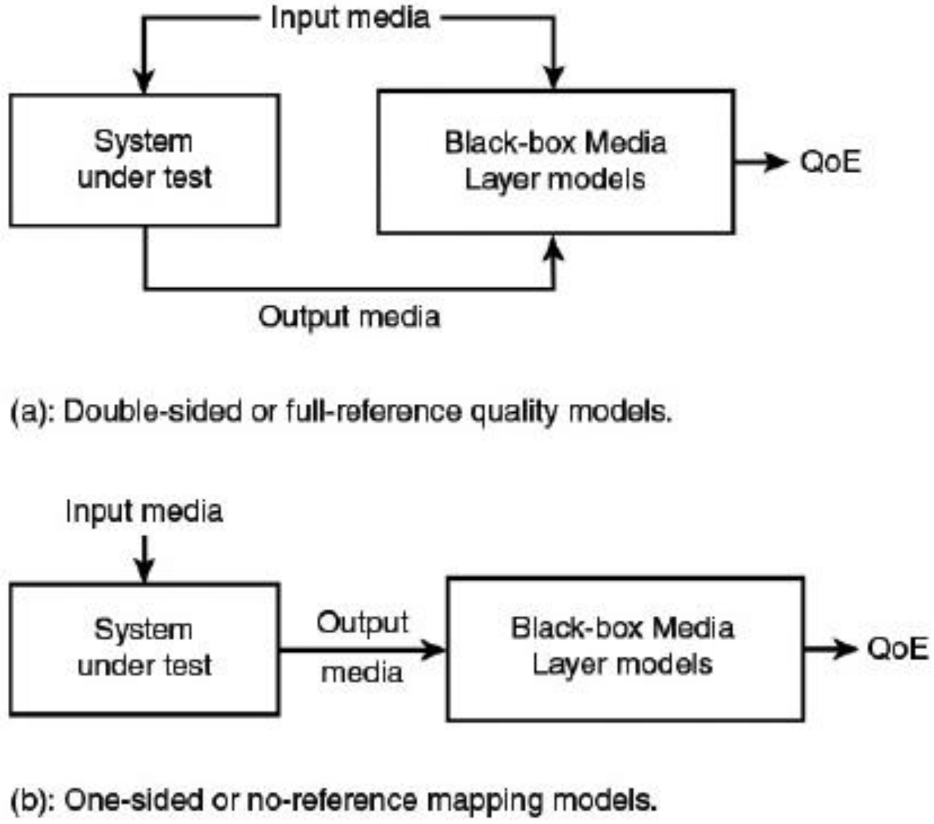

(a): Double-sided or full-reference quality models.

(b): One-sided or no-reference mapping models.

Figure 8.1 Black-Box Media-Based QoS/QoE Mapping Models

b-One-sided or no-reference quality models: They rely solely on the degraded stimulus to estimate the final QoE values. They parse the degraded stimulus to extract the observed distortions, which are dependent on the media type, for example, audio, image and video. As an example, artifacts extracted from audio stimulus include whistle, circuit noises, echoes, level saturation, clapping, interruptions, and pauses. The gathered distortions are adequately combined and transformed to compute the QoE values.

The main advantage of black-box quality models resides in their ability to measure QoE values using information gathered at the periphery of a given media processing system. Hence, they may be used in a generic fashion over different infrastructures and technologies. Moreover, it enables enhancing unconditionally quality models, that is, independently of technical and ethical constraint related to the measurement processes. Furthermore, black-box quality models may easily operate on either per-user or per-content basis.

The main shortcoming of black-box quality models resides in the requirements to access the final representation of stimulus, which is often inaccessible in practice for privacy reasons. Moreover, full-reference quality models use clean stimulus as inputs that is often unavailable or hardly accessible at the system output.

The full-reference black-box quality models are widely used for onsite benchmarking, diagnosis, and tuning of network equipment's, where clean stimulus is available. The black-box quality models are used offline for the evaluation of application-layer components, such as codec, packet loss concealment (PLC), and buffering schemes.

8.2.2 Glass-Box Parameter-Based QoS/QoE Mapping Models

The glass-box parameter-based quality models quantify the QoE of a given service through the full characterization of the underlying transport network and edge devices. The set of considered characterization parameters and their combination rules are derived based on extensive subjective experiments and thorough statistical analysis. The glass-box parameter-based models may operate off line or on line according to the availability of characterization parameters at a given measurement instant. The characterization parameters include noise, packet loss, coding scheme, one-way delay, and delay jitter. The glass-box parameter-based models are generally less accurate and coarser than black-box media-based ones.

A well-known offline glass-box parameter-based model, named E-Model, has been defined by the ITU-T in Rec. G.101. The offline glass-box parameter-based quality models are suitable for planning purposes. They enable a general overview pf QoE values of a voice transmission system at an early phase. However, for service monitoring and management, online models are needed. In such a case, the variable model parameters should be acquired at run time. This is especially suitable for IP-based services where control data, such as sequence number and time stamp, are included in each packet header. In such an environment, it is possible to extract static characterization parameters from signaling messages and variable ones from the received packets captured at the destination port. This means that parameters are acquired without acceding to the media content, which is preferable for privacy reasons.

8.2.3 Gray-Box QoS/QoE Mapping Models

The gray-box quality models combine advantages of black- and glass-box mapping models. They sample basic characterization parameters at system output in addition to some control data describing the structure of clean stimulus. The control data may be sent in separate control packets or piggybacked inside transmitted media packets. Hence, perceptually important information about a given content can be considered by the quality models. Therefore, they can measure QoE value on per-content basis. Given its simplicity to deploy and its reasonable accuracy, this class of QoS/QoE mapping models is quickly proliferating.

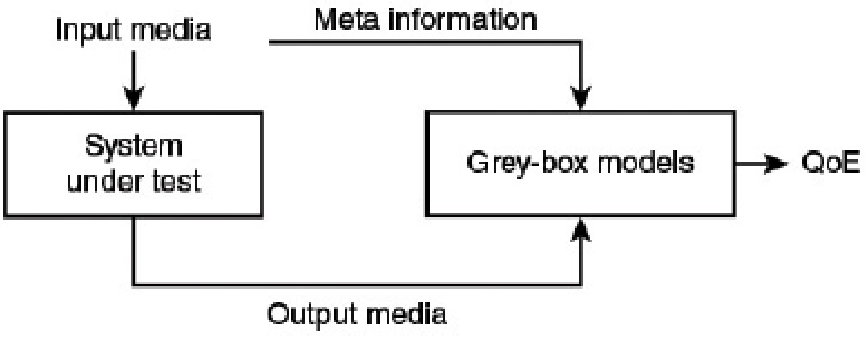

Figure 8.2 Gray-Box QoS/QoE Mapping Models

Tips for QoS/QoE Mapping Model Selection

The following checklist of five items can aid in the selection of a QoS/QoE mapping model:

- Which types of operations am I considering?
- Which parameters do I have? Can I access the signals, the contents, the packet payload or the header?
- Do I expect specifications and usage conditions to use a given mapping model?
- How much precision do I need?
- Do I have all inputs available for selected mapping models?

8.3 IP-ORIENTED PARAMETER-BASED QOS/QOE MAPPING MODELS

The area of measuring QoE of IP networks and applications is still in its infancy. However, the popularity of multimedia and user-friendly IP-based services puts QoE at the center of interest of today's ecosystem. In contrast to legacy content-oriented telecom systems (for example, public switched telephone network, radio, and TV), IP networks carry clean media content from a server to a destination using a flow of media packets composed of a header and a payload. Therefore, parameters gathered at network layer, in addition to application layer, are easily accessible at run time on user devices. This enables measuring QoE at run time using online glass- or gray-box parameter-based quality models. The QoE over IP-based networks are time- varying in contrast to telecom networks, which are roughly time-invariant. This characteristic leads to considering instantaneous and overall QoE.

8.3.1 Network Layer QoE/QoS Mapping Models for Video Services

The network layer QoS/QoE mapping models rely solely on NQoS metrics gathered from the TCP/IP stack except for the application layer (that is, transport, network, link, and physical layers). Ketyko et al. proposed the following parameter-based quality model for estimating video streaming quality in 3G environment:

$$\overline{QoE} = 8.49 - 0.02 \cdot AL - 0.01 \cdot VL - 1.12 \cdot AJ + 0.04 \cdot RSSI \qquad \textbf{(Eq. 8.1)}$$

where AL and VL refer respectively to audio and video packet loss rates, AJ and VJ represent respectively audio and video packet jitter (VJ), and RSSI is the received signal strength indicator.

Kim and Choi presented a two-stage QoE/QoS mapping model for IPTV over 3G networks. The first stage consists of combining a set of basic QoS parameters into one metric as follows:

$$\mathrm{QoS}(L, U, J, D, B) = K\{W_L \cdot L + W_U \cdot U + W_J \cdot J + W_d \cdot D + W_b \cdot B\}$$ **(Eq. 8.2)**

where L, U, J, D, and B refer, respectively, to packet loss, burst level, packet jitter, packet delay, and bandwidth. The constants K, W_I, W_u, W_J, W_d and W_b are predefined weighting coefficients, which depend on the type of the access network (that is, wired or wireless).

The second stage consists of computing QoE value as following:

$$\mathrm{QoE}(\mathrm{QoS}(X)) = Q_r\,(1 - \mathrm{QoS}(X))^{\mathrm{QoS}(X) \times A/R}$$ **(Eq. 8.3)**

where, X is a vector of parameters {L, U, J, D, B} and Q_r is a scalar limiting the range of the IPTV QoE obtained as a function of the display size/resolution of the screen. The constant A expresses the subscribed service class and R is a constant reflecting the structure of the video frames.

8.3.2 Application Layer QoE/QoS Mapping Models for Video Services

Besides NQoS parameters, application layer QoE/QoS mapping models use metrics gathered at application layers (AQoS). Moreover, they can account for the user behavior while interacting with a given video content Ma et al., the following parameter-based quality model is presented for video streaming application:

$$\mathrm{QoE} = 4023 - 0.0672\,L_x - 0.742\,(N_{QS} + N_{RE}) - 0.106\,T_{mr}$$ **(Eq. 8.4)**

where Lx refers to the start-up latency, that is, the waiting time before playing a video sequence, N_{QS} is the number of quality switches that count the number of times the video bit rate is changed during a session, N_{RE} is the number of rebuffering events, and T_{MR} is the mean rebuffering time. Khan et al., estimate QoE of a generic streamed content video over wireless networks using MPEG4 codec:

$$\mathrm{QoE}(FR, SBR, PER) = \frac{a_1 + a_2\ FR + a_3 \cdot \ln(SBR)}{1 + a_4 \cdot PER + a_5 \cdot (PER)^2}$$ **(Eq. 8.5)**

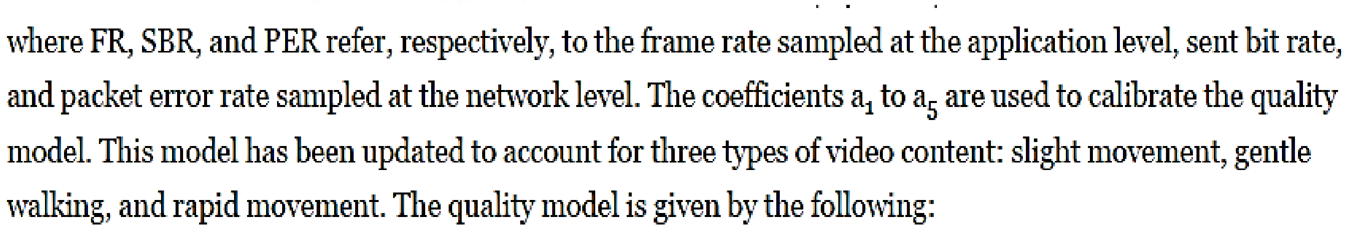

where FR, SBR, and PER refer, respectively, to the frame rate sampled at the application level, sent bit rate, and packet error rate sampled at the network level. The coefficients a_1 to a_5 are used to calibrate the quality model. This model has been updated to account for three types of video content: slight movement, gentle walking, and rapid movement. The quality model is given by the following:

$$\mathrm{QoE}(FR, SBR, BLER, CT) = a + \frac{b \cdot e^{FR} + c \cdot \ln(SBR) + CT \cdot (d + e \cdot \ln(SBR))}{1 + f \cdot (BLER) + g \cdot (BLER)^2}$$ **(Eq. 8.6)**
(Eqtn. 12.6)

where a, b, c, d, e, f, g represents constants; CT is the content type of the video; and SBR and BLER refer respectively to the sent bit rate and the bit loss error rate.

A QoE /QoS mapping model for IPTV was developed by Kuipers et al, which accounts for the startup latency and zapping time. The quality model in given by the following

$$\mathrm{QoE} = 3.5e^{-(0.15L+0.19)}N + 1.5$$ **(Eq. 8.1)**

where QoE is a one-dimension QoE component considering zapping behaviour, ZT is the zapping time expressed in seconds, and a and b are numeric constants that might be positive or negative.

8.4 ACTIONABLE QOE OVER IP-BASED NETWORKS

Actionable QoE refers to all techniques and mechanisms enabling to concretely measure and utilize QoE metrics. An actionable QoE solution strongly depends on the underlying system and services characteristics. Moreover, actionable QoE solution works over multiplane architectures that integrate data, control, and management planes. Basically, two solutions may be used to achieve actionable QoE:

1. System-oriented actionable QoE solution
2. Service-oriented actionable QoE solution

8.4.1 The System-Oriented Actionable QoE Solution

The system-oriented actionable QoE solutions account for QoE measures within the delivery infrastructure. In such a condition, services are engineered while assuming that underlying system is perfect; that is, no degradations are inserted. Figure 8.3illustrates a nominal environment where system-oriented actionable QoE solution may be provided. As can be seen, actionable QoE solution requires

- A QoS measurement module that gathers basic **key performance indicators** (KPIs) from the underlying system,
- A QoE/QoS mapping model, and
- A resource management module of controlled devices.

Each service provider specifies a target QoE level that should be offered for its customers. The QoE/QoS mapping model should be selected in a way that guarantees

(a) the availability of quality model input parameters and
(b) conformity with service specifications and conditions.

A signaling procedure may be executed to do that. The management procedure may be executed either before starting a service or during its delivery. This should be realized using an autonomous decision system, including a policy that maps observed QoE measures to a course of actions executed by managed devices.

This operational mode applies well for software-defined networking (SDN) where the network paths are managed by an SDN controller. In such a case, the measured QoE values are reported to the SDN controller, which uses them to define the behavior of SDN switches.

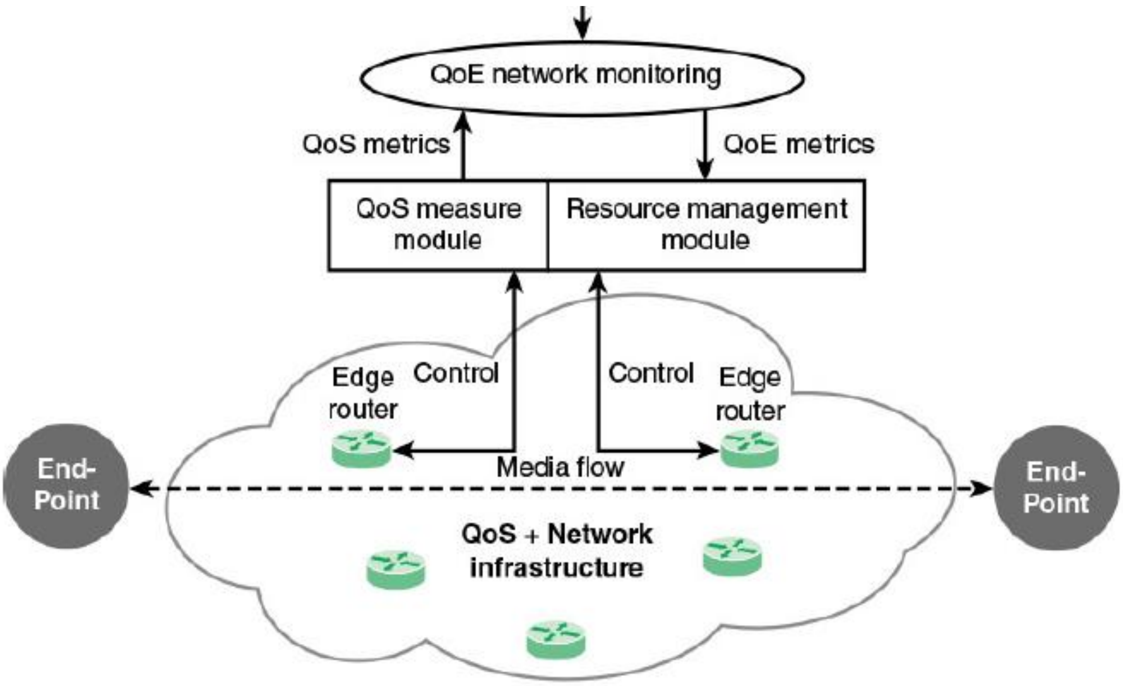

Figure 8.3 A Nominal Environment for Providing QoE-Centric Services

8.4.2 The Service-Oriented Actionable QoE Solution

The service-oriented actionable QoE solutions account for QoE values measured at endpoints and service level. In such a situation, services are engineered to deal with the underlying system flaws to reach a specified QoE level. The services may change their behavior as a function of the current context and condition. The measurement module of KPI is installed on endpoints. The QoE/QoS mapping models may be deployed either on endpoints or specialized devices. The measured QoE values are sent to endpoints to configure different application modules at sender, proxy, and receiver entities.

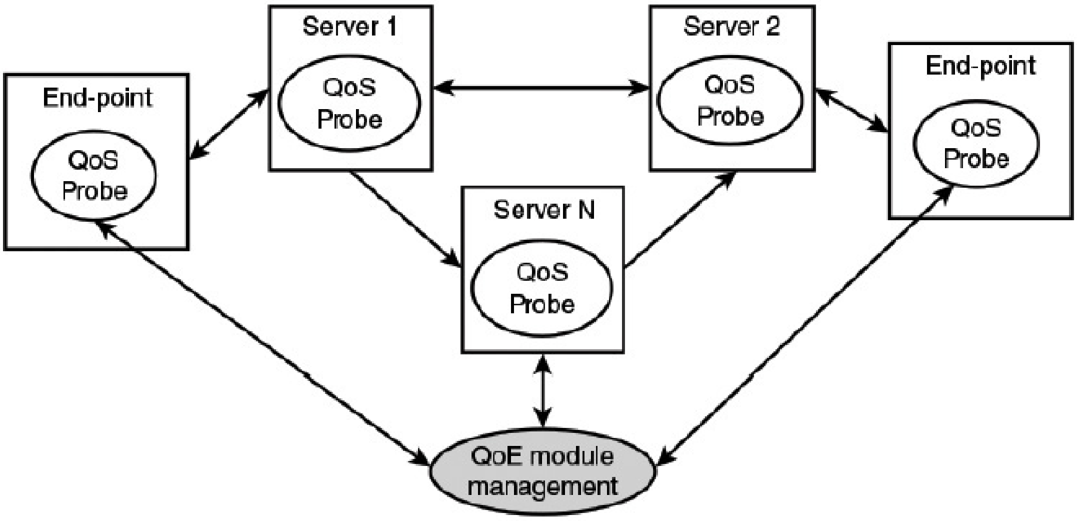

Figure 8.4 Service-Aware QoE Development Scheme

The service-oriented actionable QoE solution involves multiple advantages.

- Per-service, per-user, and per-content QoE monitoring and management solutions are performed to provide a given QoE level.

- It provides more adaptation possibility because it precisely discerns capability and the role of each service component.
- It reduces the communication overhead and balances computing loads.
- It enables component-level granularity treatment of QoE in addition to stream- and packet-level granularities.

8.5 QoE VERSUS QoS SERVICE MONITORING

8.5.1 MONITORING and its Classification

Monitoring is a strategic function that should be supported by today's IT system. It returns indicators and provides clues regarding the system performance and its workload. Moreover, it enables detecting system dysfunction and defects as well as underperforming devices and applications so that the best course of actions may be undertaken. The monitoring solutions of current IT systems may be classified into the following four categories:

- **Network monitoring:** Provides measures about performance of paths and links used to deliver media units. They are collected at packet processing devices (router and switch) and may operate on per-flow or per- packet bases. The path characterization metrics, such as throughput, packet loss, reorder and duplication, delay, and jitter, are calculated using atomic metrics extracted from packet header, such as sequence number and time stamps.
- **Infrastructure monitoring:** Provides measures about devices performance and resources state, such as memory, CPU, IO, load, and so on.
- **Platform monitoring:** Provides performance indicators about the computing center where back-end servers are running. They may work over a virtualized infrastructure where business application logics are deployed using virtual machines.
- **Service monitoring:** Provides measures about services performance. The metrics are dependent on each application, and may be realized from technical or perceptual perspectives.

Typically, the monitoring solution in a distributed system involves a variety of **probes** that measure the performance of a given element participating in the service delivery chain. It includes a reliable and scalable manager that remotely configures probe behavior, especially in terms of the frequency and content of measurement reports.

The monitoring solution should provide communication facility between the manager and the managed devices. The interaction between the manager and managed entities is conventionally realized using the connectionless User Datagram Protocol (UDP) through a couple of reserved ports.

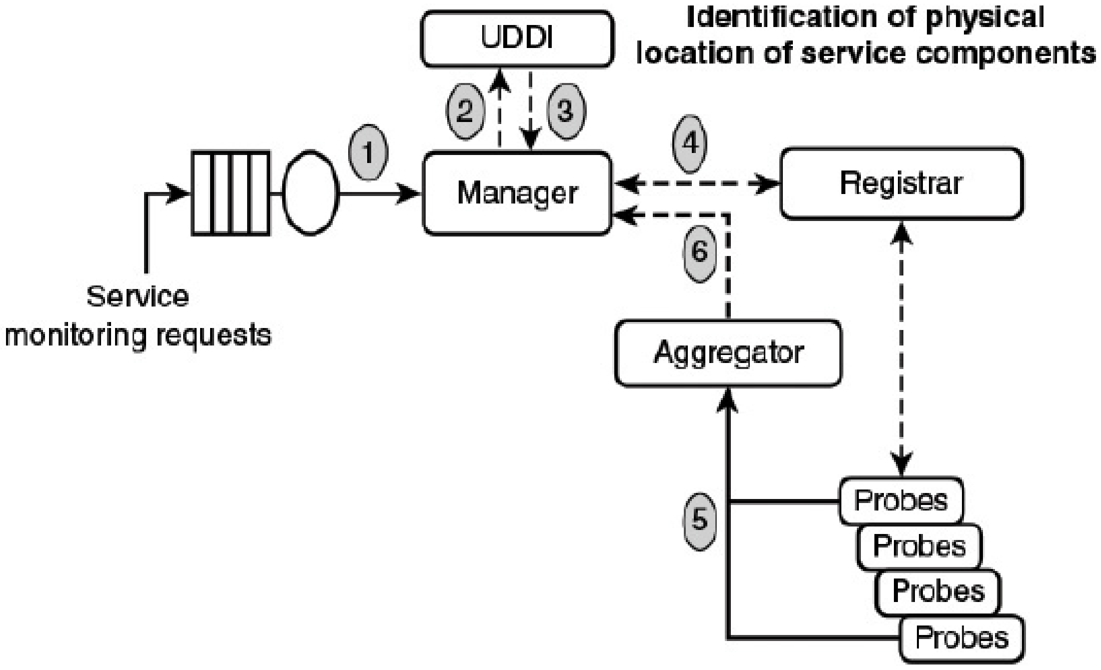
UDDI
Identification of physical location of service components
2
3
1
4
Manager
Registrar
6
Service monitoring requests
Aggregator
5
Probes
Probes
Probes
Probes

FIGURE 12.6 A Baseline and Generic Monitoring Solution

Figure 8.5 A Baseline and Generic Monitoring Solution

Figure 8.5 presents a typical configuration of an on-demand monitoring solution. The manager receives monitoring requests from customers expressed using specific syntax. Upon the receipt of a new monitoring request, the manager inquires with a Universal Description, Discovery, and Integration (UDDI) directory to get more information about the monitored service, such as location and properties. The monitoring solution, including a set of probes, is deployed over a given infrastructure. They register themselves automatically once a monitored component is activated in a preconfigured registrar. The registrar keeps traces and features of all active probes. They should be configured off line by the administrator to report metrics according to a given behavior during a service. The metrics generated by the probes may be aggregated and processed before sending them to the manager that performs data analytic procedure.

8.5.2 QoS Monitoring Solutions

The emerging QoS monitoring solutions are basically developed for data centers and clouds where virtualization technology is supported. Figure 8.6 shows a network- and infrastructure-level monitoring solution built for cloud-based IPTV service. The audiovisual content servers are placed on a cloud. The traffic sent from the content servers to IPTV devices is permanently monitored through a set of Vprobes deployed across the network. A Vprobe is an open-ended investigatory tool that is used in the cloud environment to inspect, record, and compute the state of the hypervisor as well as each virtual machine running service business logics. The flows of video packets are parsed at different measurement points. The information collected by Vprobes is used next to reconstruct service-level detailed records (SDRs). Each record contains the most relevant information of the complete session between an origin (server) and a destination (user). The critical parameters of the messages associated with an IPTV session are stored inside the SDRs.

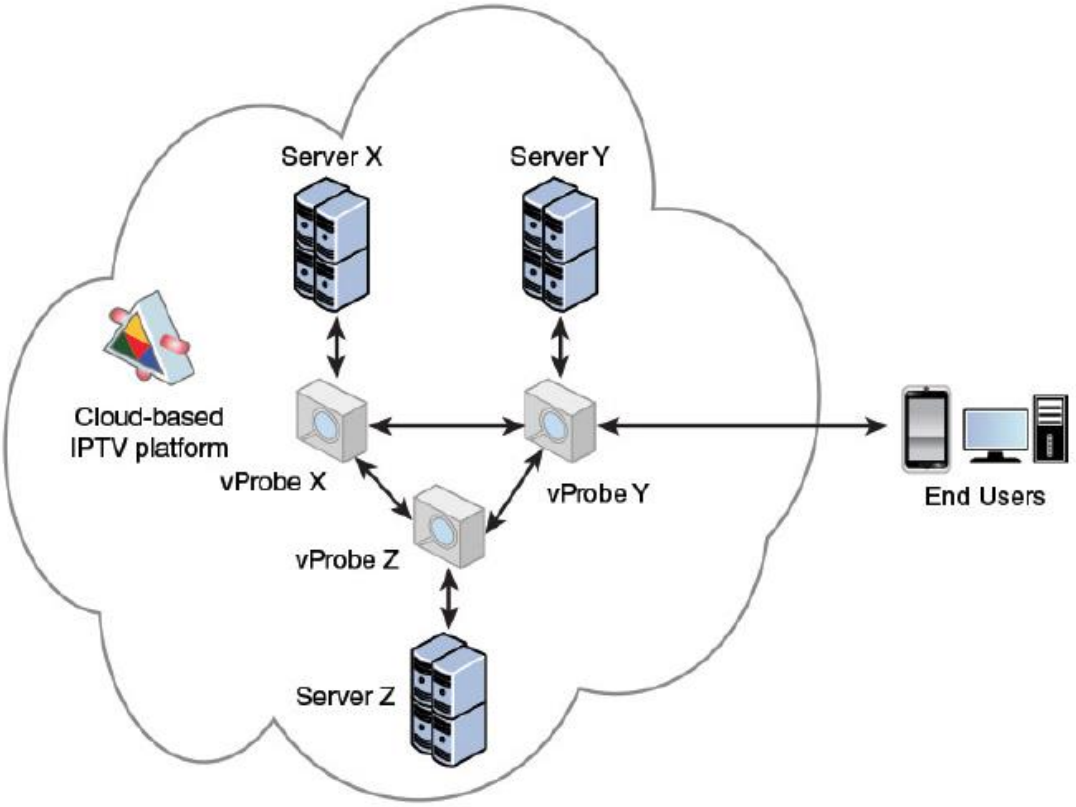

Figure 8.6 vProbes Approach in Cloud-Based IPTV Network

8.5.3 QoE Monitoring Solutions

The QoE monitoring solution strongly depends on QoE/QoS mapping models. The diagrams in Figure 8.1 show four configurations that can be used to monitor at run time the QoE values of IP-based video streaming services. The configurations differ in term of the measurement and mapping model locations. Each configuration is denoted using XY expression, where X refers to the measurement location and Y refers to the quality model location. They may take one of these values: N for network, C for client, and B for both.

A. **Static operation mode (NN):** Both the measurements of KPIs and QoE are performed inside the network. The QoE/QoS mapping model is installed on a device listening to the service delivery path. The quality model uses collected KPIs, prior knowledge about video coding schemes, and endpoint characteristics. The characteristics of endpoints can be acquired either by polling them or by inspecting exchanged SDP (Session Description Protocol) messages. The QoE measurement points may include an endpoint emulator enabling a realistic reconstruction of received streams.

B. **Nonembedded dynamic operation (BN):** The measurement of KPIs is performed at both the network and the client, whereas QoE values are measured inside the network. The quality model uses gathered KPIs, prior knowledge about coding schemes, and information about the client obtained using a customized signaling protocols.

C. **Nonembedded distributed operation (CN):** The measurement of KPIs is performed at the client side, and these are sent periodically to the QoE/QoS mapping model located inside the network.

D. **Operation embedded (CC):** The measurement of KPIs and QoE is performed at the client side. The QoS/QoE mapping model is embedded inside the client. The measured QoE metrics may be reported to a centralized monitoring entity.

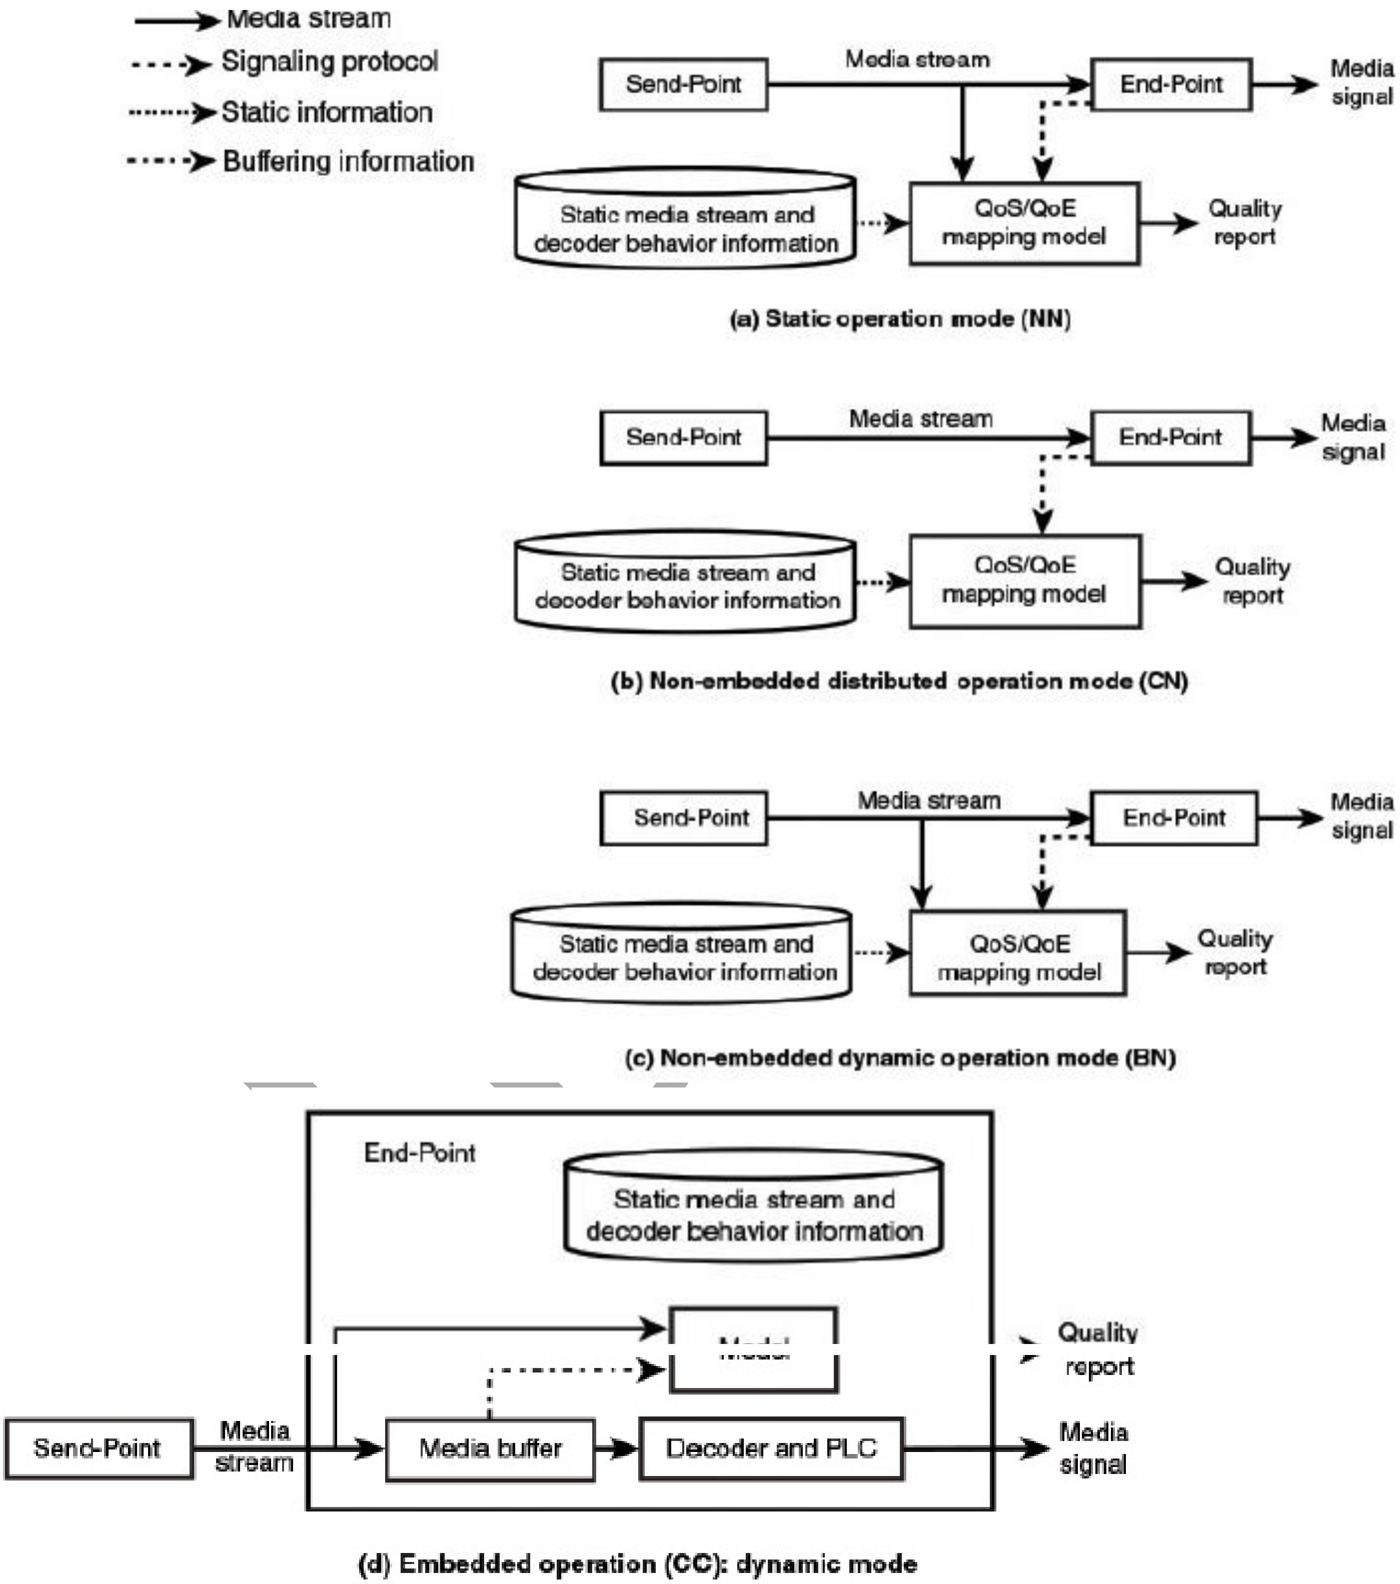

(a) Static operation mode (NN)

(b) Non-embedded distributed operation mode (CN)

(c) Non-embedded dynamic operation mode (BN)

(d) Embedded operation (CC): dynamic mode

Figure 8.1 The Operational Working Modes of Quality Models in Networks

The architecture of QoE-agent is based on a layered definition of APIs that enable convenient grouping of different factors that influence QoE. The six layers are defined as follows:

- **Resource:** Composed of dimensions representing the characteristics and performance of the technical system(s) and network resources used to deliver the service. Examples of such factors include network QoS in terms of delay, jitter, loss, error rate, and throughput. Furthermore, system resources such as server processing capabilities and end user device capabilities are included.
- **Application:** Composed of dimensions representing application/service configuration factors. Examples of such factors include media encoding, resolution, sample rate, frame rate, buffer sizes, SNR, etc.
- **Interface:** Represents the physical equipment and interface through which the user is interacting with the application (type of device, screen size, mouse, etc.).
- **Context:** Related to the physical context (e. g. geographical aspects, ambient light and noise, time of the day), the usage context (e.g. mobility/no-mobility or stress/no-stress), and the economic context (e.g. the cost that a user is paying for a service).
- **Human:** Represents all factors related to the perceptual characteristics of users (e.g. sensitivity to audiovisual stimulus, perception of durations, etc.).
- **User:** Users' factors that are not represented in the Human layer. These factors encompass all aspects of humans as users of services or applications (e.g., history and social characteristics, motivation, expectation, and level of expertise).

8.6 QOE-BASED NETWORK AND SERVICE MANAGEMENT

The quantified QoE values may be considered in networks and services management. This enables getting an optimal trade-off that maximizes QoE and minimizes consumption of resources. The major challenge resides in the translation of QoE metrics into a course of actions that definitely enhance encountered QoE and reduce resources consumptions.

8.6.1 QoE-Based Management of VoIP Calls

The management of Voice over IP (VoIP) based on QoE has been extensively investigated in the literature. The goal is to maintain a constant QoE level during a whole packet voice session transmitted over time-varying quality IP networks. Typically, QoE measurement probes following one glass-box parameter-based model are installed on VoIP endpoints. They collect at run time atomic KPIs, which are transformed and given as inputs to a QoE/QoS mapping model. After a new measure of QoE values is received, a QoS controller adjusts the reconfigurable network parameters within a delivery path, such as queuing allocation and congestion thresholds.

8.6.2 QoE-Based Host-Centric Vertical Handover

Mobile consumers over next-generation networks could be served at one moment by several overlapping heterogeneous wireless networks. The network selection/switching procedure can be performed either at the start or during the service. An internetwork hard handover occurs when users switch from one network to another because of specific reasons related to both consumers and providers. A handover could be managed in network- or host-centric way. In a host-centric approach, however, end nodes can perform a handover when quality of service becomes unsteady and unsatisfactory.

Figure 8.8 illustrates a likely envisaged scenario where the client could be served either by WiMAX or Wi-Fi systems. Appropriate equipment should be deployed and configured, such as outdoor and indoor units, server, router, and Wi-Fi and WiMAX access points to enable network handover. Throughout a vocal call, the client may switch from WiMAX system to Wi-Fi system, and vice versa.

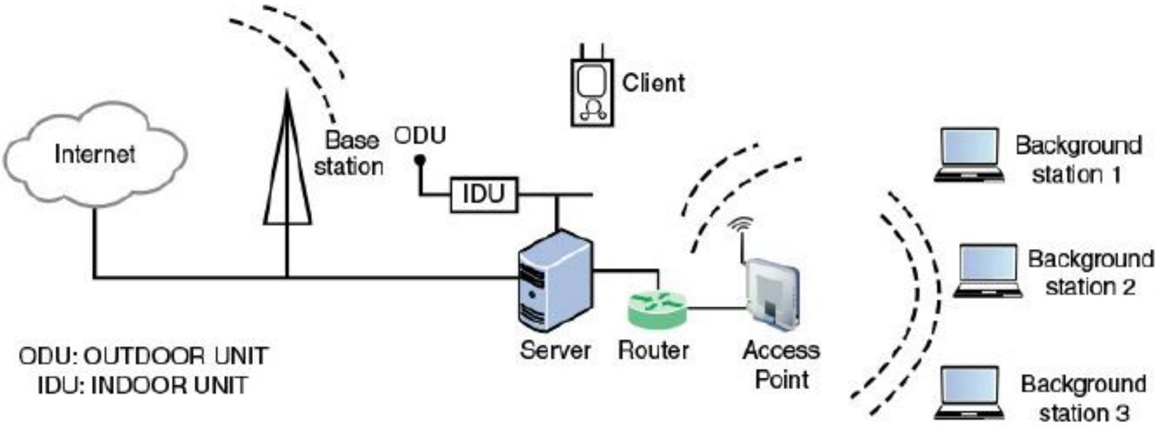

Figure 8.8 Network Selection Wi-Fi and WiMAX Based on Client and Link Quality

8.6.3 QoE-Based Network-Centric Vertical Handover

The goal is to perform a handover between overlapping WLAN and GSM networks. This allows, on one hand, relatively exploiting the high capacity of a WLAN, and on the other hand, reducing the GSM network load and cost. Figure 8.9 shows a scenario where a mobile subscriber initiates a voice call to a landline PSTN subscriber using a WLAN as a last-wireless hop. Next, when the QoE of the voice call goes below a given critical threshold because of mobility or congestion, a handover is performed. In such a case, the mobile subscriber is linked to the landline subscriber using the GSM infrastructure. The hands-free terminal is equipped with two wireless card interfaces to allow connection to WLAN and GSM networks. The mobile terminal sends adequate "quality reports" to a PBX that analyzes received feedbacks. Once an unsatisfied score is detected, the PBX instructs the mobile terminal to perform a handover. To do that seamlessly, a voice channel is opened using GSM infrastructure between the mobile terminal and PBX, which is responsible to relay received voice information toward the fixed subscriber.

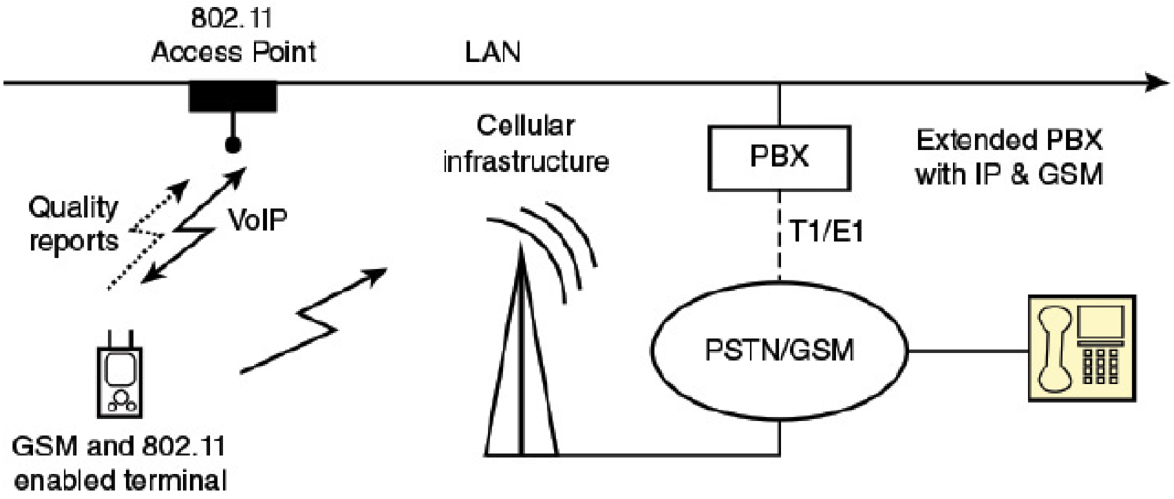

Figure 8.9 Handover Scenario Between WLAN and GSM Networks