ctc technology & energy

engineering & business consulting

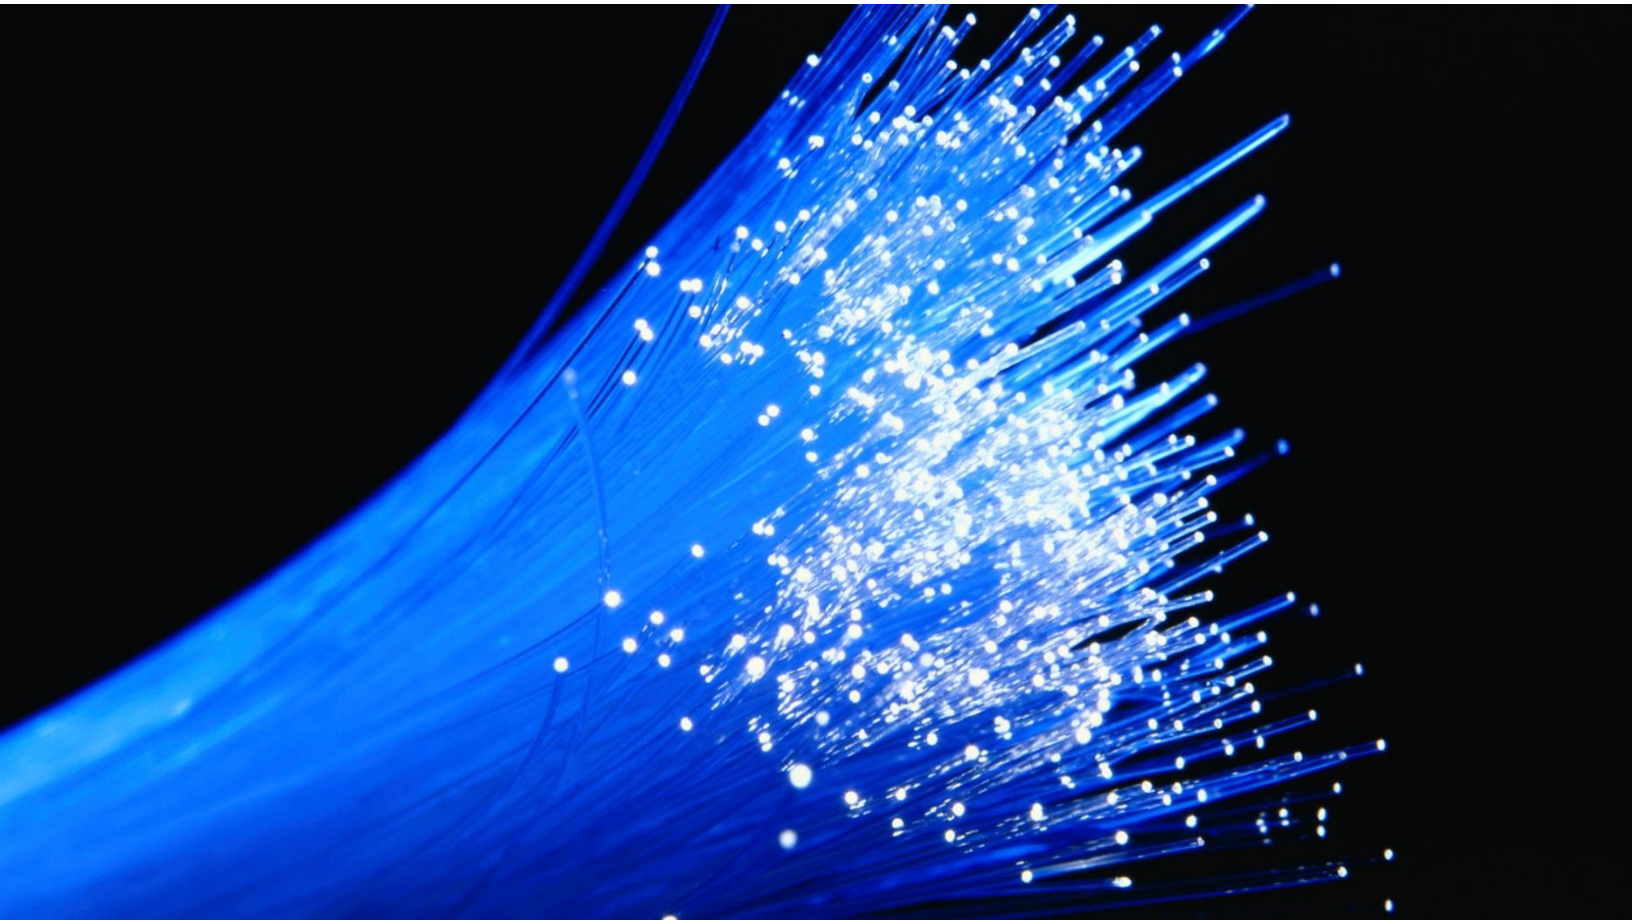

Technical Guidance Regarding Broadband Infrastructure for Libraries

Prepared for the John S. and James L. Knight Foundation

FINAL | October 2017

Columbia Telecommunications Corporation

10613 Concord Street • Kensington, MD 20895 • Tel: 301-933-1488 • Fax: 301-933-3340 • www.ctcnet.us

Contents

Figures

Tables

1 Introduction

Public libraries have long been places where anyone can go to access knowledge—and in that tradition, high-speed internet access has become a primary feature of today's libraries. Beyond the migration of traditional library services to the internet, some patrons have come to rely on libraries as the primary, or sometimes only, place where they can get online to do homework, check email, apply for jobs, and communicate with family and friends. Many users are also looking for high-bandwidth connections with greater capacity than they can access at home or elsewhere in the community.

Though many libraries report that they have upgraded their broadband service in recent years, nearly one-third report their bandwidth rarely or only sometimes meets demand.[1] Those libraries also report that their bandwidth barely meets the recommendations set out by the American Library Association and the Federal Communications Commission (FCC).[2] This is an issue not just for providing services to patrons but for meeting the day-to-day needs of library staff.

Libraries need technical guidance to make informed decisions about how to improve broadband connectivity and create robust networks. Most community libraries simply do not have sufficient staff resources or technical expertise to secure adequate and competitively priced high-speed connections and network services for their institutions. This was the common refrain we heard while researching this report—including in conversations with leaders from the American Library Association; the Universal Service Administrative Company (USAC), which administers the federal subsidies under E-rate; and libraries across the United States in a wide range of environments.

This report provides tools and guidelines that libraries can use to evaluate their current and future needs and find broadband services that fit their budgets and goals. Using this report, a library or library system will be able to specify needs such as bandwidth, quality of service, and network type. This will enable the library to select potential solutions and, if appropriate, create an effective request for proposals (RFP) document and evaluate the benefits and risks of solutions that a service provider may propose.

This report was prepared by CTC Technology & Energy under a grant from Knight Foundation. CTC is grateful to Knight Foundation for supporting our educational mission.

[1] "2014 Digital Inclusion Survey: Survey Findings and Results Extended Summary," Information Policy & Access Center, University of Maryland, October 1, 2015, https://goo.gl/IAB4Rf.

[2] 100 Mbps download speeds at libraries that serve fewer than 50,000 people and 1 Gbps (1,000 Mbps) at libraries that serve more than 50,000 people. In our companion paper for the Knight Foundation, "The Library of the Future and Its Need for Bandwidth," we address the bandwidth demand that libraries will face in the future, as well as the opportunities for libraries to use this bandwidth to expand their missions through new and innovative services. "Comments of the American Library Association," GN Docket No. 16-245, Sept. 2015, https://goo.gl/bggP8P.

2 Understanding Broadband Technologies

To make informed decisions about connectivity services, libraries need to understand their options. This section provides a high-level overview of the connectivity technologies in use today, as well as considerations for each technology.

2.1 Fiber Optics

Fiber is the most advanced and highest-capacity form of wireline communications. It is used to connect most medium- to large-sized institutions. In addition, fiber is built to wireless cell sites and access points, and to cabinets and nodes on copper and cable modem networks. In almost all cases, fiber is nearby even if a library's direct connection is through one of the other media discussed below.

Since the 1980s, fiber has been incorporated into middle-mile and backhaul connections—the lines that are used to aggregate data traffic and provide high-capacity transport between cities and across continents. Fiber optic cables contain numerous thin strands of glass (or in some cases plastic) that carry data as a series of pulses of light, traveling from one end of the fiber to the other. A backbone fiber cable could have hundreds of strands. A fiber cable serving a neighborhood or a few buildings would have a few dozen strands. A cable to an individual apartment or house might have one or two strands.

Optical signals can travel great distances with minimal signal deterioration. Typical fiber networks can carry broadband data signals up to 50 miles between electronics. This is a major change from the electrical signals of metal conductor-based telephone and cable television networks. Fiber cables and their optical signals do not experience most of the physical limitations of metal-based networks. The superior range eliminates the need for electrical power and equipment in the middle of most networks. And fiber networks have lower operating costs than cable and DSL networks, because they require less staffing and maintenance.

Fiber networks are also more reliable than networks built with other technologies such as coaxial cable and copper. With less equipment needed to operate the network, there are fewer points of failure that could disrupt communications. Optical fibers do not conduct electricity and are immune to electromagnetic interference. These properties allow optical fibers to be deployed where conductive materials would be dangerous or ineffective, such as near power lines or within electric substations. Further, fiber optic cables do not corrode due to weather and environmental conditions in the same way that metallic components can deteriorate over time.

Once installed, fiber optic cables have few technical limitations. The main drawbacks for fiber optic networks are the up-front cost and the process of building out to connect new locations where fiber does not yet exist. The price for fiber optic cable itself is declining, but costs associated with construction to existing premises remain high. (However, in "greenfield" areas, the cost of fiber optic network construction is the same on a per-unit basis as coaxial or DSL.)

Fiber networks can be continually upgraded to faster speeds. Fiber provides a broad communications spectrum and has a capacity of thousands of gigabits per second (Gbps) per individual fiber with off-the-shelf networking hardware. Even lower-priced equipment easily provides 1 Gbps service. The main limitation on the speeds fiber networks can achieve is not based on the properties of the fiber optic cables themselves but on the processing power of the networking equipment connected to the network. Fiber's ability to scale has led some to describe it as "future-proof."

Fiber networks using "Active Ethernet" or comparable technologies provide symmetrical download and upload speeds—in contrast to DSL or cable broadband services, which typically provide faster download than upload speeds. Networks that also have fast upload speeds are useful for institutions and businesses because they facilitate the sharing of extremely large data files. For example, libraries connected by fiber may be able to share digital resources like video and audio as quickly as if they were stored locally. Fiber networks can scale to meet the demands of the next generation of internet services and applications without any additional construction.

2.2 Service Delivery over Fiber

Fiber connections available to libraries fall into four general categories:

- Lit fiber services
- Dark fiber lease
- Dark fiber IRU
- Self-provisioned (library-built) fiber

2.2.1 Overview

"Lit fiber" refers to fiber strands that are actively being used; the strands are connected to equipment that transmits light (containing the data) over the strands. In the context of internet services, lit fiber services are those where the internet service provider (ISP) owns and operates both the fiber and the equipment connected to it at the library and at the ISP's facilities.

Lit fiber is always a service—the library connects its equipment to the service provider's fiber or equipment and the service provider manages the service beyond the demarcation point, essentially as a "black box." The library typically knows little or nothing about the fiber route or intermediate devices on the network. This arrangement can be desirable if the library does not have the expertise or desire to maintain or operate fiber, and if the lit fiber arrangement is projected to be cheaper over the lifetime of the service and to perform better than an arrangement where the library lights and maintains the fiber and electronics itself.

A variety of services can be offered over lit fiber; these include Dedicated Internet Access (DIA) and wide area network (WAN) services (discussed below). Fiber may provide multiple services to multiple customers.

In contrast, "dark fiber" refers to fiber strands that are not actively being used. In the context of internet services, dark fiber also refers to strands that are provided or constructed in a dark state for the library itself to connect and light. While a library receiving lit fiber is always receiving a *service*, a library that owns or operates dark fiber is operating *infrastructure*.

In a dark fiber scenario, a library may lease or buy strands within a service provider's cables, or it may construct its own fiber cables and place them on poles or underground. The library may manage the network itself or it may contract with another organization. In a dark fiber scenario, the fiber and equipment used by the library network are not shared by other customers on the provider's network.

When considering a choice between lit and dark fiber, a useful analogy may be leased and owned real estate. Like leased property, lit fiber services may be costlier than dark fiber options over long periods of time because the library pays a recurring fee to the service provider for management and maintenance of the network. Additionally, if the service provider needs to build fiber to connect a library because no fiber yet exists at the location, the service provider may charge the library an up-front construction or connection fee in addition to monthly charges, or build that cost into the monthly charges. This cost will vary depending on the provider's business model and the library's proximity to existing fiber infrastructure.

The library is also locked into the service offerings (speeds, technology) of the service provider and is highly dependent on the quality of maintenance and management performed by the service provider. The service provider operates all routers and switches and connections to the internet, and the library's only recourse in the event of problems is through its contract.

Despite the higher long-term cost and other potential drawbacks, a lit fiber service may be a good option for libraries that do not have multiple locations, that cannot or do not wish to pay the up-front cost of dark fiber, or that cannot or do not wish to operate the network.

2.2.2 Lit Fiber Services

The services offered over a fiber connection can vary greatly depending on the technology used by the service provider and by the needs of each library branch or system. These services may allow control over quality of service (including prioritization of different types of network traffic), multiple distinct data paths, and the ability to quickly add new locations without major changes or interruptions on the network. (See Appendix B for more details.)

2.2.3 Dark Fiber Lease

A dark fiber lease gives the library exclusive use of fiber strands from one library to another for the duration of the lease. Typically, a library-to-library connection requires one or two strands. The service provider leasing the fiber brings the fiber inside the library and terminates it on a panel, where the library connects it to its electronics. Leasing agreements are typically for a period of a few years, and recurring fees are generally paid on a monthly or yearly basis.

A lease is generally not as cost-effective over the long term as an IRU agreement (described below), but this option may be necessary if the library cannot afford the up-front payment required by an IRU. The service provider performs maintenance and repair on the fiber, but the library has full end-to-end control of the fiber strands. In contrast to a library using a lit service, a library that leases fiber will own and control all electronics connected to the network, and those electronics will determine the speed of the connection and the particular network protocol being used (e.g., Ethernet, SONET). The library will be able to upgrade this service simply by upgrading the electronics, independent of the service provider.

2.2.4 Dark Fiber IRU

An Indefeasible Right of Use (IRU) agreement is effectively a long-term lease that grants the library exclusive use of one or more strands of dark fiber. IRU agreements typically last 20 years. The cost of the IRU is paid up-front with an additional maintenance fee usually paid annually. For customers who are able to afford the up-front payment and have this option available, IRUs provide significant cost savings compared with leased fiber. From the technical perspective, a dark fiber IRU is the same as a dark fiber lease.

2.2.5 Self-Provisioned (Library-Owned) Fiber

A self-provisioned fiber network is built, owned, and maintained by the library. In this scenario, the library builds fiber from location to location. It places an entire fiber cable, which can contain hundreds of fiber strands. This provides spare capacity for the library, which will only need a small percentage of the strands. This strategy also makes it possible to build the fiber in partnership with other entities—often the library's home municipality or county—that use the remainder of the capacity for other public purposes such as serving local government needs. Joint planning with other public sector entities is always advisable. Another option is for the fiber to be built as part of a public–private partnership, with the other fiber cables built and used by fiber and wireless service providers.

By building and maintaining its own fiber network, a library system will have the greatest flexibility and control over the network. It will not have to pay fiber lease or IRU fees. On the other hand, this option will place the greatest demand on the library in terms of staffing and training because the library will be responsible not only for operating and maintaining the network electronics at each location but also for maintaining the physical fiber network itself.

The library will need to negotiate attachment agreements with utility pole owners and/or construct underground fiber in the right-of-way. This approach has been used by governments and community anchor institutions for years and has become increasingly popular as the bandwidth demands of those institutions increase. The FCC made self-provisioned networks eligible for E-rate funding in 2016 so long as a library shows it to be the most cost-effective option.

2.3 Other Network Media

Other types of broadband connections include hybrid fiber-coaxial (cable TV), digital subscriber line (copper telephone), wireless technologies, and mobile broadband services.

2.3.1 Hybrid Fiber-Coaxial (HFC)

Cable broadband technology is currently the primary means of providing broadband services to homes and small and medium-sized businesses in urban, suburban, and small-town areas in the United States. Coaxial cables were originally designed to provide video services. One legacy of this design is the highly asymmetrical capacity of coaxial systems, with 95 percent of the physical capacity originally designed for the downstream direction. In order to provide sufficient capacity for data services and reduce the noise in long runs of coaxial cable, cable operators built fiber to each neighborhood to within a mile or so of each customer, typically lashing the fiber to the coaxial cables. Thus, coaxial cable networks have transformed into hybrid fiber-coaxial (HFC) networks.

Because of HFC's relative ubiquity in non-rural areas and its inherently greater capacity than commercial wireless service and copper telephone lines (the medium underlying digital subscriber line, or DSL, service), HFC cable networks will continue as the dominant broadband communications technology for most homes and businesses for the foreseeable future.

2.3.1.1 Technical Capacity and Limitations

Although there are a number of significant limitations inherent in cable systems relative to fully fiber optic networks, cable system capabilities will increase over the next few years with the deployment of new technologies and the extension of fiber closer to customers.[3]

In an HFC network, headend or hub locations house the core transmission equipment. Fiber connections extend from these hubs to multiple nodes, each of which serves a given geographical

[3] Cable is not as scalable "out of the box" as communications systems that were designed from the outset to provide Internet-type broadband data services. Issues include coaxial cable's limitations in terms of physical capacity, a physical architecture optimized for broadcast communications, and a significant remaining migration path to full end-to-end internet protocol (IP) operations.

area (e.g., a neighborhood). These optical nodes are electronic devices located outdoors, attached to aerial utility lines, or placed in pedestals. The equipment in the node converts the optical signals carried on fiber into electronic signals carried over coaxial cables. Coaxial cable then carries the video, data, and telephony services to individual customer locations. Figure 1 illustrates an HFC network.

Cable operators have extended fiber optics progressively closer to their subscribers but, for cost reasons, have generally stopped at nodes about one mile from the premises. Comcast, for example, typically only constructs fiber to the premises of customers that subscribe to Metro Ethernet and other advanced services.

The current leading cable technology for broadband data, known as data over cable service interface specifications version 3.0, or DOCSIS 3.0, makes it possible for cable operators to increase capacity by bonding multiple channels together. The DOCSIS 3.0 standard requires that cable modems bond at least four channels, for connection speeds of up to 200 Mbps downstream and 108 Mbps upstream (assuming use of four channels in each direction). A cable operator can carry more capacity by bonding more channels.

It is critical to note that these are peak speeds and that the capacity is shared by all customers—typically hundreds of homes or businesses—on a particular segment of coaxial cable. Speeds may decrease during bandwidth "rush hours," when more users simultaneously use greater amounts of bandwidth. For example, residential bandwidth use typically goes up considerably during evening hours, when more people use streaming video services and other large data applications.

Figure 1: DOCSIS 3.0 Network Architecture

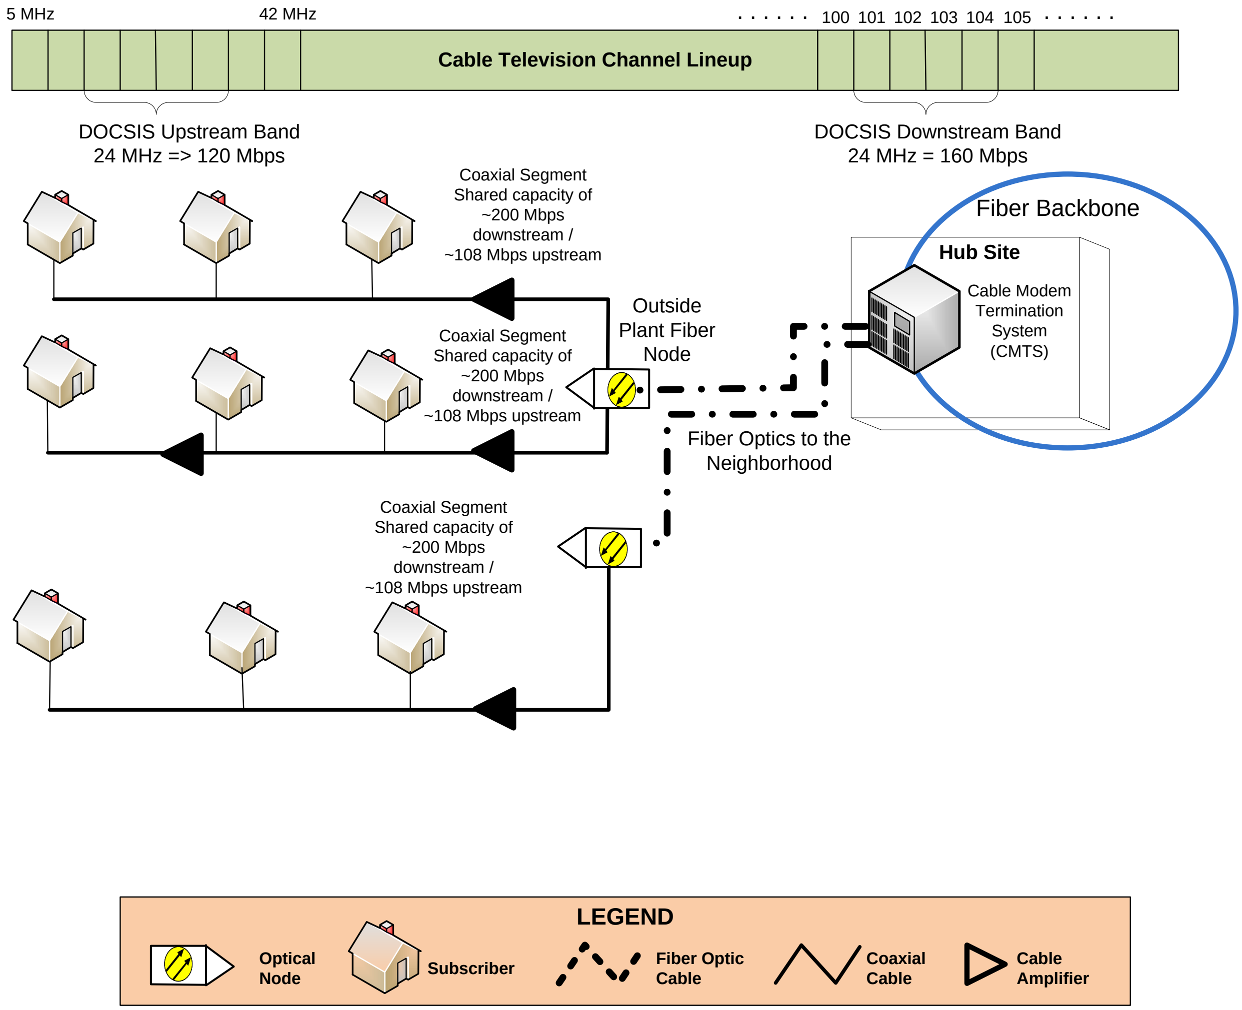

Although the standard is still in the test phase, both Comcast and Time Warner have announced plans to begin upgrading their systems to DOCSIS 3.1. The cable industry claims that DOCSIS 3.1 will eventually provide 10 Gbps downstream capacity and 1 Gbps upstream. This will not be possible for most actual cable systems. Even with high-efficiency implementations, the actual downstream capacity for a shared node area will be closer to 6 Gbps, and that capacity will be shared by a few hundred users.

2.3.1.2 Factors Impacting Quality and Speed of Service

The following factors will determine a cable broadband customer's service speed and quality:

1. **Bandwidth capacity of cable plant.** Most coaxial portions of a cable network have a capacity of 750 or 860 MHz, but they can be upgraded to 1 GHz and beyond. If the cable corrodes, the available bandwidth shrinks, limiting possible connection speed.

2. **Number of customers sharing a fiber node.** Cable capacity is shared among all the users connected to a given fiber node, so connection speeds may decrease significantly during peak usage hours. Cable companies can reduce the number of customers sharing a node by putting fiber deeper into their systems and moving the node closer to the customers. Cable operators manage capacity by monitoring usage and upgrading on a neighborhood-by-neighborhood basis.

3. **Proximity of customer to fiber.** Another advantage of moving fiber closer to the customer is that signals travel shorter distances on coaxial cable. With progressively shorter stretches of coaxial cable, the inherent problems with reliability and interference decrease.

4. **Standards and protocols.** Cable operators can make faster connection speeds available by dedicating more channels to data services and upgrading their networks to later versions of industry standards. DOCSIS 3.1 makes more efficient use of available spectrum, freeing up more bandwidth for data download and upload.

Cable operators often offer services with “blast” or “burst” speeds of “up to” more than 100 Mbps. Although a customer may be able to access these speeds on occasion, the actual speeds available may be significantly lower during peak usage hours.

In the early days of cable internet service, some cable companies provided a free cable modem to schools and libraries, so cable service became the first internet connection for many libraries. Over time, many libraries continued that arrangement, sometimes upgrading to faster business-grade services over cable modems. However, cable modem service is highly asymmetric and now not fast enough for many libraries. As a result, libraries that obtain service from a cable operator are often actually obtaining a lit fiber service.

2.3.2 Digital Subscriber Line (DSL)

During the last century, phone companies connected nearly every home and business in the United States to a copper telephone wire. Copper has a fraction of the bandwidth capacity of coaxial cable and suffers from greater signal loss and interference. But because of copper’s ubiquity, digital subscriber line (DSL) technology over copper has been an important way for people and libraries to connect to the internet.

While DSL has been a useful retrofit of existing infrastructure, copper cable is reaching its physical limitations as a broadband medium and will not be able to meet future bandwidth needs of U.S. libraries.

The main determinant of maximum DSL speed for a given copper telephone line is the length of the copper line from the telephone company central office. In systems operated by large telecommunications companies, the average length is 10,000 feet, corresponding to DSL speeds between 1.5 Mbps and 6 Mbps. In systems operated by small companies in rural areas, the average length is 20,000 feet, corresponding to maximum speeds below 1.5 Mbps.

The fastest copper telephone line technologies widely deployed in the United States are VDSL and VDSL-2, the technologies underlying AT&T's U-verse and comparable services. Because these technologies use high frequencies, they are limited to 3,000 feet over typical copper lines—a much shorter distance than in most HFC systems—and require fiber to the node (FTTN). Therefore, in order to operate VDSL and VDSL-2, telecommunications companies must invest in large-scale fiber optic construction and install remote cabinets in each neighborhood.

In practice, telephone companies using VDSL-2 over highly upgraded copper lines have been able to provide 25 Mbps over a single copper pair and 45 Mbps over two pairs to the home or business—but it took a significant investment to make it possible for a small percentage of the copper phone lines to temporarily keep pace with cable. Providing even greater speeds will require some combination of deeper fiber construction, a breakthrough in transmission technology over copper lines, and conditioning and upgrading of the existing copper lines.

The Nokia (Alcatel-Lucent) G.Fast DSL product reportedly reaches speeds of 500 to 800 Mbps in various environments—but it only delivers such speeds for a few hundred feet, which essentially requires the construction of fiber to the curb in front of each home or business—an investment that would be comparable to building a full fiber network.[4] As a result, G.Fast has so far mostly been focused on deployments using telephone wires already installed inside older office or apartment buildings.

Rural libraries and small libraries are the most likely to be connected using DSL. The lines carrying the DSL service may be the same copper telephone lines that were used for the original voice lines and T1 data connections. However, DSL technology is not able to support increasingly symmetrical speeds in the hundreds of Mbps, and most libraries that make use of it will need to move to other technologies.

2.3.3 Wireless

Wireless technologies can provide a solution in low-density rural areas where the high cost of building wired networks often leaves rural residents without a wired broadband option. Wireless internet service providers (WISPs) are potentially able to fill these coverage gaps, sending signals from base stations to antennas on or near customer premises. WISPs are not able to offer connection speeds on a market-wide basis comparable to cable or fiber built to each premises.

[4] See, for example, "Nokia and Frontier Communications deploy G.fast technology to expand gigabit ultra-broadband access across Connecticut," News Release, Nokia, May 25, 2017, https://www.nokia.com/en_int/news/releases/2017/05/25/nokia-and-frontier-communications-deploy-gfast-technology-to-expand-gigabit-ultra-broadband-access-across-connecticut.

However, they may be the best available solution if cable or fiber is not cost-effective. Even in an urban setting, a WISP can create a point-to-point network from rooftop to rooftop with individual links in the Gbps range, connecting to fiber where it exists.

2.3.3.1 Technical Capacity and Limitations

Fixed wireless networks built with off-the-shelf equipment today tend to have an aggregate capacity between 100 and 250 Mbps. With help from the latest antenna and signal processing technologies—innovations like higher-order multiple input, multiple output (MIMO) antennas, and the use of spatial multiplexing—these capacities will likely increase to as fast as 750 Mbps. Smaller WISPs use the same unlicensed spectrum bands as Wi-Fi, which does not have strong long-distance transmission qualities.

Most wireless networking solutions require the antenna at the library to be in the line of sight of the base station antenna. This can be especially challenging in hilly or mountainous regions. It is also a problem in areas with dense vegetation or multiple tall buildings. WISPs often need to lease space at or near the tops of radio towers; even then, some customers may be unreachable without the use of additional repeaters. And because the signal is being sent through the air, climate conditions like rain and fog can impact the quality of service.

Some wireless providers in rural areas have begun to use vacant television frequencies called TV white space (or simply white space) to provide service. These TV bands have much better non-line-of-sight transmission qualities than the unlicensed bands; however, because white space technology is still in an early phase of development, compatible equipment is far more expensive than other off-the-shelf wireless equipment. However, it is possible that costs could come down rapidly in the coming years—a development that would make white space an attractive solution for rural libraries.

Meanwhile, wireless equipment vendors offer a variety of point-to-multipoint and point-to-point solutions. A library would likely need a point-to-point solution with dedicated bandwidth from the service provider to obtain sufficient bandwidth and quality, and even then the upstream capacity would be limited. (Small businesses and residences can get by with a point-to-multipoint solution.)

2.3.3.2 Factors Impacting Quality and Speed of Service

The following factors will determine a fixed wireless customer's service speed and quality:

- **Wireless equipment used.** Different wireless equipment has different aggregate bandwidth capacity and uses a range of different spectrum bands, each with its own unique transmission capabilities.

- **Backhaul connection.** If a WISP cannot get an adequate connection back to the internet from its tower, wireless equipment upgrades will not increase available speeds beyond a certain point.

- **Unobstructed line of sight.** Most wireless networking equipment requires a clear, or nearly clear, line of sight between antennas for optimum performance.

- **Weather conditions and foliage.** Depending on the spectrum used, weather conditions like rain or fog may cause interference. Also, line-of-sight paths that are clear during the winter may be obstructed by foliage during the warmer months.

- **Future capacity and lifespan of investment.** Wireless equipment generally requires replacement every five to 10 years, both because exposure to the elements causes deterioration and because the technology continues to advance at a rapid pace, making decade-old equipment mostly obsolete. The cost of deploying a wireless network is generally much lower than deploying a wireline network, but the wireless network will require more regular investment.

As with fiber, a library may "self-provision" a wireless network, installing antennas on rooftops and connecting multiple links back to the internet. In this case, the library would need to perform the engineering and installation, obtain space on towers and rooftops, and perform needed maintenance.

2.3.4 Mobile Broadband

Cellular wireless carriers have been consistently increasing their data speeds with the rollout of faster and higher-capacity technologies, such as Long-Term Evolution (LTE).[5] Over the past few years, they have begun providing data plans with speeds comparable to and in many cases greater than a typical residential customer's internet service.

Wireless providers operate a mixture of third-generation (3G) and fourth-generation (4G) technologies, typically selling devices (telephones, smartphones, air cards, tablet computers) bundled with 3G or 4G services.

The strict definition of 4G from the International Telecommunications Union (ITU) was originally limited to networks capable of peak speeds of 100 Mbps to 1+ Gbps, depending on the user environment; according to that definition, 4G technologies are not yet deployed. In practice, however, a number of existing technologies (e.g., LTE) are called 4G and represent not only a speed increase over 3G technologies but also a difference in architecture—more like a data cloud

[5] LTE is a 4G cellular wireless technology offering data speeds of typically around 30 Mbps.

than a cellular telephone network overlaid with data services. The ITU and other expert groups have more or less accepted this.

Because of its limited bandwidth and the cost of the bandwidth (typically metered), mobile broadband is not well suited to connecting libraries. However, mobile devices are useful tools for library patrons because of their ubiquity and the ease of connecting to free library Wi-Fi connections. Indeed, the proliferation of mobile devices and patrons' reliance on them underscores the need for robust library networks. And some libraries, viewing their missions of inclusion and education in the broadest terms, even lend patrons mobile hotspot devices that give the patrons access to mobile broadband when not at the library.[6]

Current 4G technologies will give way in many markets to "5G" technologies, which, where and when they are deployed, starting in 2020, would provide higher speeds and lower latency (meaning, a shorter time delay) than 4G. However, 5G is also expected to use frequencies that have a far shorter effective range, requiring denser and more numerous small transmitters. We do not anticipate that 5G will ever be superior to the underlying wireline networks. And like current wireless technologies, 5G will rely on fiber backhaul.

5G technologies are expected to begin appearing in 2020 and to include a mixture of higher speed mobile technologies using high-frequency "millimeter wave" spectrum and upgraded versions of 4G. Given that 5G is still in the early stages, several issues are not yet clear:

1. **The applications that are driving the technology.** 4G was driven by smartphones and the migration of high-speed data to mobile use; 5G may be driven by more demanding data use and by machine-to-machine and automated vehicle use.

2. **The type of technology and spectrum, and the delivery of service.** The extent to which 5G will be deployed as enormous numbers of small antennas on utility poles and lights, and the degree to which it will be deployed in outside high-density urban and suburban corridors remain to be seen.

3. **The business case.** Unless there is a business need that justifies the cost of the upgrade, whether and where service providers will spend the money and political capital to undertake a massive 5G deployment is unknown.

---

[6] "Wi-Fi Hot Spots for Rent: How Public Libraries Are Changing with the Times," http://www.govtech.com/network/Wi-Fi-Hot-Spots-for-Rent-How-Public-Libraries-Are-Changing-with-the-Times.html (accessed March 2017).

2.3.5 How Libraries Can Leverage 5G

There are three main ways libraries might participate in this expansion and the eventual transition to 5G: (1) by facilitating the deployment of antennas on library property, given the likelihood that wireless carriers will need to install many more small antennas with 5G, and indeed are already doing so with 4G; (2) by providing wireless carriers with a way to connect their antennas back to the internet; and potentially (3) by exploring ways that 5G technology might enable new applications for library patrons and their communities.

Libraries in heavily trafficked areas or that are themselves heavily trafficked may be attractive locations for the new antennas that are often called "small cells" due to their short range. Libraries can offer a way for wireless carriers to attach to existing structures, place new poles, or replace existing structures with poles that serve the same function but are designed to host wireless attachments as well. These types of installations may also require an easement to allow the wireless provider to run fiber on library property to the small cell.

Libraries that have built their own networks may also be able to provide wireless carriers with a way to connect their small cells back to the internet or to another carrier site by leasing spare conduit, dark fiber, or lit fiber services. Libraries that are planning to build a fiber network may also be able to reduce construction costs by cooperating with wireless carriers, which may need to install fiber in the area.

And finally, libraries might potentially be able to serve their communities by demonstrating or hosting new applications that leverage 5G, which, as a White House announcement last year put it, "will enable breakthrough applications for consumers, smart cities, and the Internet of Things that cannot even be imagined today."[7]

[7] "Fact Sheet: Administration Announces an Advanced Wireless Research Initiative, Building on President's Legacy of Forward-Leaning Broadband Policy," White House, July 15, 2016, https://obamawhitehouse.archives.gov/the-press-office/2016/07/15/fact-sheet-administration-announces-advanced-wireless-research.

3 Assessing Current and Future Needs

To plan effectively, a library must first understand how its patrons and staff connect to the internet and other networks, whether its current services and connections are sufficient, and how growing demand and new projects and services will impact its needs in the future.

We recommend against looking at the current connection speeds and services and simply assessing them as “OK,” “a little slow,” or “way too slow.” In our experience, this approach can sharply underestimate both current and future needs.

First, if the connection is already too slow, patrons and staff may have “given up” on the internet available at the library, assuming from previous experience that it is too slow for most uses and electing not to use the library internet for video or other more data-intensive tasks. For other users, it is not always clear whether slow network speeds are caused by the internet connection or a component within the library, such as cabling, network electronics, or a poorly maintained public computer.

It is better to use a bottom-up, holistic approach that considers the number of users, the likely applications, and the current utilization metrics—and then builds in margin for future growth. This report suggests a methodology for this type of approach for a wide range of libraries and library systems.

The key questions in this approach are:

- How many users are on the network?
- What is the type and amount of usage?
- What is the flow of the data?

The number of users and amount of usage will vary over the course of a typical day and week and over longer periods of time. The flow of data will mostly be to and from the internet. But it is important to know whether there are other library locations that need to be interconnected, whether there are ways to centralize the internet connection, and whether it is possible to take advantage of regional networks or connection points to the internet.

Through the process described in the following sections, a library or library system can develop technical parameters to describe both the network it currently needs and the network it is likely to need in the future. Expressing these parameters in standard network terms will enable a library or library system to characterize its needs for a service provider, and will minimize uncertainty and questions during the process of improving its service.

Requirements should include, in rough order of importance, bandwidth, symmetry, quality of service (QoS), topology, and customer service response. These are defined below.

3.1 Bandwidth

For purposes of this report, bandwidth is the capacity of the connection or connections from the library to the internet, or from the library to the outside network. Bandwidth is measured in bits per second, with a "bit" being either a 1 or 0, the smallest unit in a data stream.[8]

For reference, a typical connection bandwidth for businesses and institutions (including libraries) in the 1990s and 2000s was 1.5 Mbps (millions of bits per second), corresponding to a "T1" circuit usually carried on a copper telephone line. Connection bandwidth evolved as demand increased, to faster speeds of tens of Mbps, carried over advanced "DSL" telephone lines, cable modem systems, point-to-point wireless systems, or fiber optics.

In recent years, the norm for larger institutions has become hundreds or thousands of Mbps (1,000 Mbps equals one gigabit per second, or Gbps). The typical approach to delivering Gbps speeds is with fiber optics, though in some circumstances, gigabit speeds can be delivered with point-to-point wireless technologies, and the cable industry is currently developing means of delivering these speeds over a mixture of fiber and coaxial cable.

It is important to note that bandwidth needs will dictate the type of physical connection that can be employed. As the need increases to higher bandwidth, most service providers will opt to use fiber to deliver the service, owing both to its high bandwidth capacity and the ability of fiber optic cable to easily operate at higher and higher bandwidth simply through upgrading the electronics at the end of the network (i.e., without laying new cable).

Given that upgrading a network connection may require a different physical medium, it is important to plan ahead. In some cases, a network upgrade will require new cables to be constructed to the library, or providing the connection may be beyond the capability of the service provider.

Fiber optics and other physical network types are described in more detail in Section 2.

3.2 Symmetry

A "symmetrical" internet connection offers the same level of bandwidth going downstream (network to library) as upstream (library to network). An "asymmetrical" network has an imbalance; usually with download speeds much faster than upload speeds. Different network technologies provide different levels of symmetry, and different institutions have different needs for symmetry. Consumer-grade network connections tend to be highly asymmetrical, with a

[8] For example, a character in text is coded as a series of eight ones or zeros. The letter B is coded in standard ASCII code as "01000010." So to send a 140-character tweet you need 1,120 bits.

factor of 10 or more times the speed in the downstream direction compared with the upstream direction.

This is both due to the demand from the typical user in the past (more focused on web browsing and viewing content than on creating content and originating video streams) and due to the technology, which has mostly been focused on demand and is significantly cheaper to provide in an asymmetrical format.

Libraries at which the networks are mostly used by patrons who are downloading data have download-to-upload ratios of about 10:1. But libraries at which the networks are centered around cloud applications and administrative use, and whose patrons are using interactive video and other two-way applications, often have ratios closer to 1:1.

Figure 2 illustrates an asymmetrical network with a downstream/upstream ratio of 5:1 (top) and a symmetrical network (bottom).

Figure 2: Network Connections Can Be Asymmetrical or Symmetrical

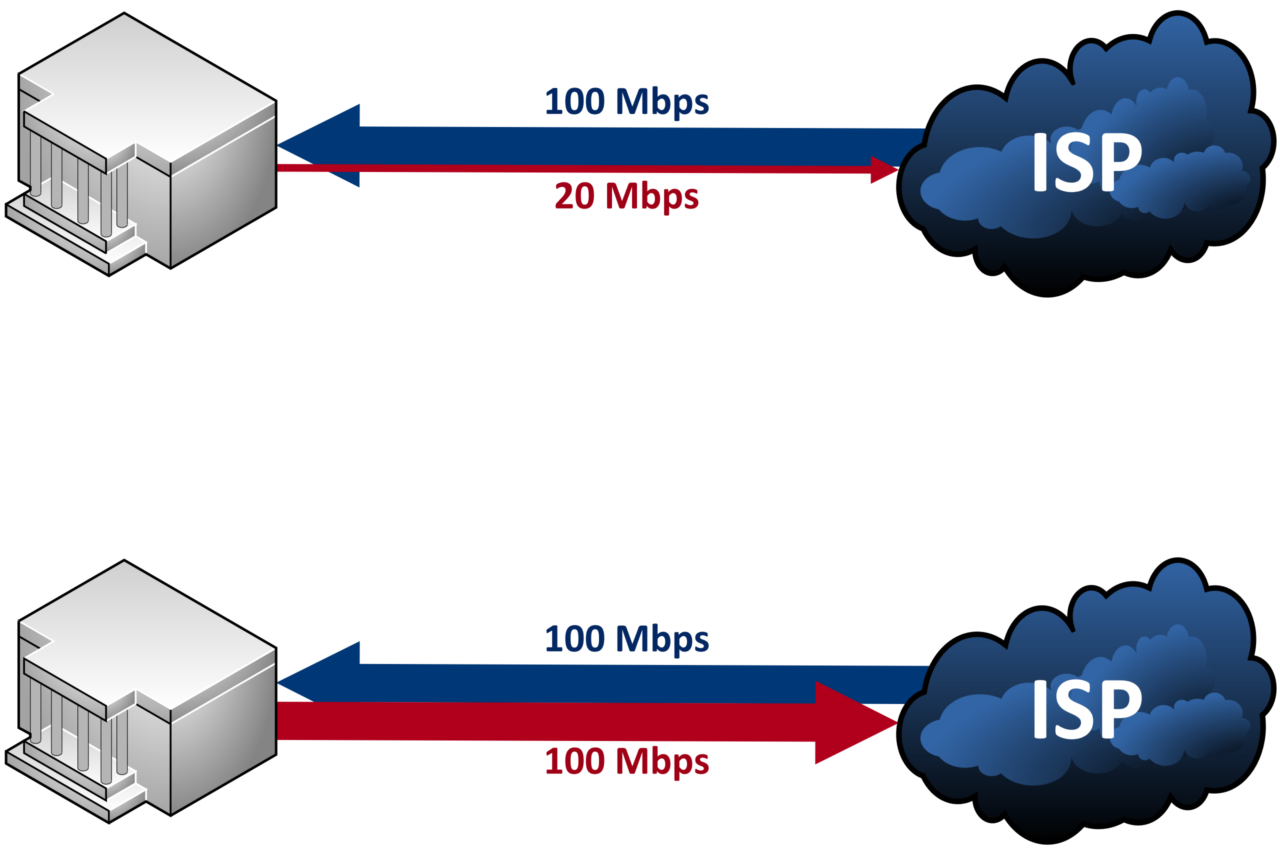

3.3 Quality of Service (QoS)

QoS refers to the availability, continuity, and consistency of a communications service. Some applications, such as voice calls and video conferencing, are especially sensitive to QoS.

QoS is measured by network equipment and devices. The metrics include the percent downtime, the time delay over the network, the variation in that delay, the severity of network problems, and the time spent to respond to and fix a problem.

Given that QoS is the product of many separate connections and components, it varies from provider to provider and from service to service. It also varies based on the utilization of the network, both by the library itself and by any other users, businesses, and institutions also served on a network. The specification needs to include sufficient guarantees so that a network provider consistently delivers a service that meets the library's data needs and does not fail to deliver for certain applications or times of the day.

Network service providers frequently oversubscribe a network connection. For example, a connection to a library may be a certain speed (100 Mbps, for example), but it may be connected at a central office to a switch with several other 100 Mbps connections, which are in turn connected over a 100 Mbps connection to an internet backbone. In many cases, such as residential service, this can work well and provide significant economies. However, if not adequately managed, oversubscription may result in the network not consistently delivering the committed bandwidth, especially during peak internet usage times.

The terms of QoS are stated in a service level agreement (SLA). Most residential and small business users do not have an SLA, or have an SLA with no quantitative, enforceable metrics. In the industry, this is referred to as "best effort," meaning the provider commits to doing its best but that the outcome is not guaranteed. Some SLAs have guidelines for performance but no penalties, while others have very mild penalties.

A library determines the QoS it requires based on its needs, including the types and criticality of network communications at each branch. For example: In the event of a network outage, what types of library activities would be interrupted or canceled? Under what circumstances would the library have to close? Are there voice and video needs that require especially trouble-free connections? The library may need to ask a service provider the variation in cost of service under different SLA terms to inform its decisions.

It is not reasonable to expect that a network connection will be working 100 percent of the time, nor is it cost-effective to build it that way. Thus, it is important to determine the amount of downtime that the library is willing to accept. This can be done by reviewing which aspects of the library's day-to-day operations, such as patron activities, access to the resource catalog, or the ability to take payments, would be affected by an outage. It is also important to specify whether downtime outside the library's normal operating hours or during scheduled maintenance outages is counted.

QoS parameters include uptime, latency, jitter, and loss.

Uptime (or availability) is a measurement of how much time a network connection is functioning. This metric is usually expressed as a percentage of time over the course of a year. For example, a network connection with 99.9 percent uptime would have been down for approximately eight hours and 45 minutes over the course of a year. This metric is sometimes referred to by the number of nines in the percentage, so a 99.99 percent uptime would be called "four nines."

Latency (or delay) measures how long it takes for data to get from one computer to another, jitter measures the variation in that delay, and loss measures how much data is sent but never reaches its destination.

Like downtime, some amount of delay, jitter, and loss will be present on most network connections. In general, applications can compensate for these problems, but too much of any of them can still cause noticeable performance disruptions.

If a network is up but is not working adequately (e.g., performing slowly or failing to comply with other QoS parameters) the network can still be considered out of compliance with availability standards, or considered partially available or impaired, depending on how the SLA is structured.

3.4 Topology

"Topology" refers to both the physical layout and the logical function of the network. The physical topology describes the location of devices, cables, and connections. The logical topology describes how data travels through the network.

Different topologies can be used to accomplish the same goals but may vary in terms of cost, performance, and features. For the most basic topology, with one library and one connection to the internet, it is relatively straightforward to assess whether the library's needs will be met by the connection. A more complex topology might have multiple libraries connected by an internal library network that connects to a central point that connects to the internet.

If a library depends on an intermediate network for connectivity to the internet, the internet connection must be able to accommodate the total amount of bandwidth required by all libraries. For example, if two libraries each require 1 Gbps of bandwidth and share one internet connection, the shared connection must be able to accommodate at least 2 Gbps.

Likewise, any requirements for delay, jitter, and loss must be met at all points in the network to ensure that no quality problems are introduced in the connections between the libraries. If a library system anticipates traffic between its own library branches for digital resources, email, or any other services hosted at a central library, then connections between the libraries must accommodate that traffic in addition to internet traffic. For this reason, it may be necessary to have internal connections that are faster than the internet connection.

Figure 3 through Figure 5 illustrate three example topologies: libraries with independent network connections, libraries with network connections aggregated through an internet connection to a central library, and a library mesh network—that is, a network where any library branch can connect with any other without requiring an internet connection to do so—aggregated through an internet connection to the central library.

Figure 3: Branch and Main Libraries with Independent Network Connections

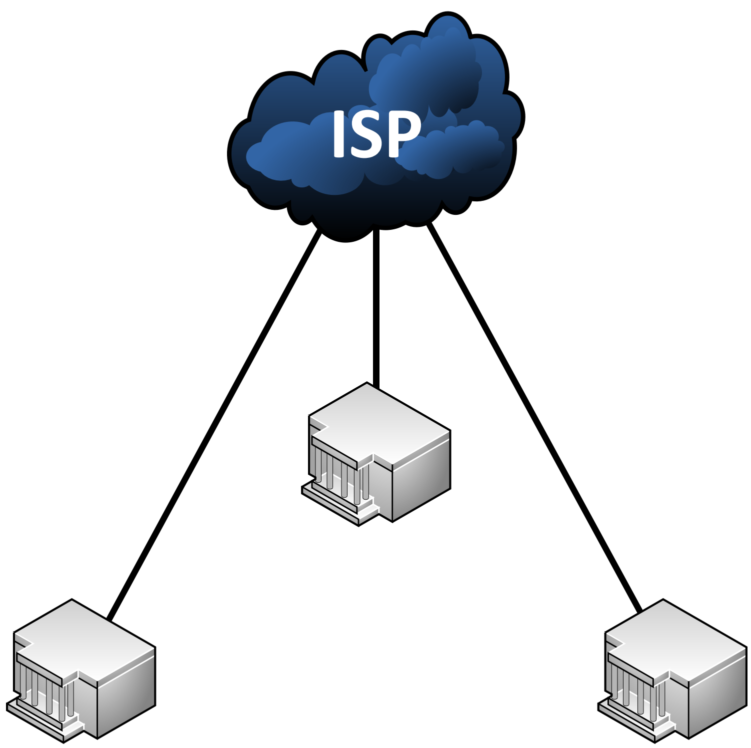

Figure 4: Branch Libraries Aggregated Through Central Library

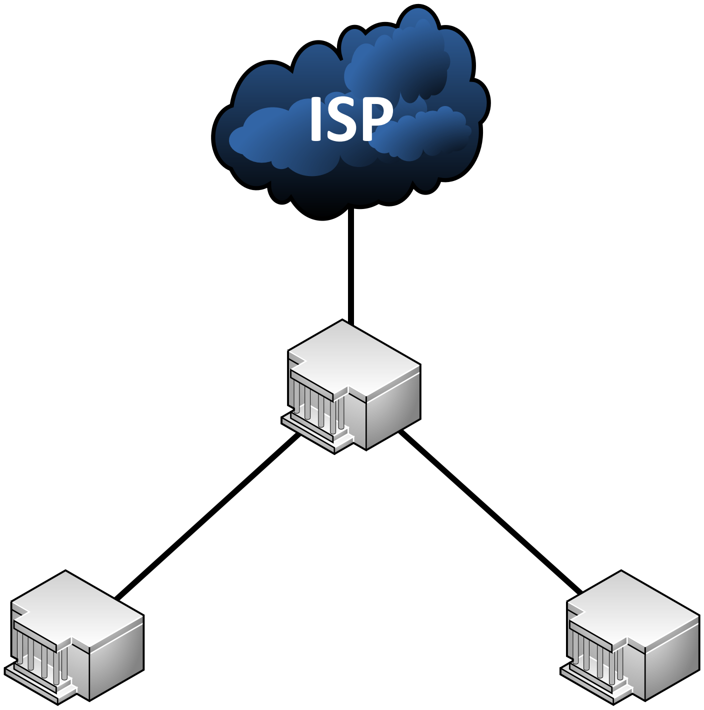

Figure 5: Library Mesh Network Connected Through Central Library

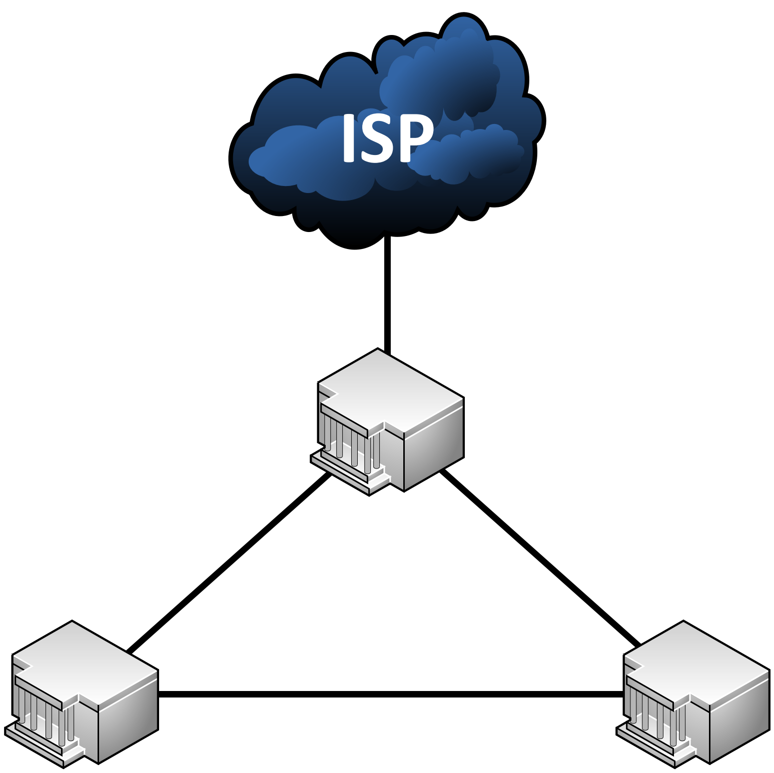

3.5 Customer Service

When selecting an internet service provider, libraries must take customer service into account, using criteria beyond bandwidth and technical numbers. Long response times and uncertainty about repair of network services can undermine the entire library's mission. A service provider also needs to be well matched to the technical team managing a library network—not only available but also skilled in responding to the specific needs of a library and having procedures and work hours that match the library's needs.

Customer service criteria should include prompt availability in the event of a problem, as well as response time for working on the problem, on site if necessary. Criteria should also include time for restoration of the service, sometimes with temporary before permanent fixes. Finally, a service provider needs to describe a contact and escalation procedure, preventive maintenance procedures, monitoring practices, real-time and historical reporting data, and characteristics of the help desk and network operations center. If there is an SLA, the provider needs to describe how it monitors the SLA and how penalties, if any, are determined.

4 Measuring Network Performance

As anyone who has had trouble with a home internet connection can attest, purchasing a fast connection does not guarantee that the network will always be fast. There may be times when web pages are slow to load, when streaming video looks choppy, or a VoIP call sounds garbled even though the available bandwidth should be sufficient for those tasks. To the person trying to use the network, the conclusion is the same: “The network is slow.”

In addition to bandwidth, internet service providers (ISPs) generally measure the quality of service (QoS) parameters of uptime, delay, jitter, and loss discussed in Section 3.3. We discuss strategies libraries can employ to measure these QoS parameters below.

4.1 Web-Based Speed Tests

There are many online tests available to help you estimate the speed of your network connection. These include:

http://beta.speedtest.net/
https://www.measurementlab.net/tests/
https://speedof.me/

Your service provider may host its own bandwidth test as well. These tools provide a snapshot view of available bandwidth and can be used to estimate the performance of the network as experienced by a user. These tests are not perfect, however, and several caveats should be kept in mind:

1. **Results can be affected by other traffic on the network.** If other users are downloading files or streaming video while you run a test, the amount of bandwidth available to you will be reduced and the test results will be lower. Try to run the test outside of business hours while no bandwidth-intensive tasks (such as online backups or software updates) are running.

2. **The LAN connection may be slower than the internet connection.** Bandwidth tests are specific to the computer you are using. If your computer has a slow wired or wireless connection to your network, that will impact the results. Make sure to use a computer with a 1 Gbps wired network interface connected to a 1 Gbps network port. Connect the computer as close to the library’s router as possible (with few devices between the computer and the internet connection). If your internet connection is faster than 1 Gbps, it may not be possible to use a web-based speed test.

3. **Results can be affected by the location of the speed test server.** Just like your local network, the internet is made up of many different connections. A speed test server that

is connected directly to, or run by, your service provider, may yield faster results than a server on a remote network. Network congestion or heavy use of the speed test server may also yield slower results. Running multiple tests at different times of the day may help give you a better idea of overall performance.

4. **Results will often be lower than the theoretical maximum speed.** Because there is a certain amount of overhead involved in network traffic, speed test results will likely be lower than the physical speed of the connection. On the other hand, if your internet service is burstable—meaning that you are allowed to temporarily exceed the amount of bandwidth you pay for—results may actually be higher than expected. The service provider should be able to provide some guidance about what to expect.

5. **Results might be measured in bits or bytes.** Network speeds are typically measured in bits, while the size of a computer file is typically measured in bytes. A byte is made up of 8 bits and is typically denoted with a capital B, whereas a bit is typically represented by a lower-case b. Bandwidth tests often offer measurements in both bits and bytes, so make sure you know which is being used. Bits are typically measured in kilobits per second (Kbps or Kb/s) or megabits per second (Mbps or Mb/s). Bytes are typically measured in kilobytes per second (KBps or KB/s) or megabytes per second (MBps or MB/s). 100 Mbps is equivalent to 12.5 MBps.

4.2 Bandwidth Monitoring

Another approach to monitoring performance is to collect and graph bandwidth information from the router that connects to the internet. This approach shows you how much bandwidth the library is actually using, as opposed to merely telling you what data speed is being delivered to any one device at a given moment. The range of options for network monitoring software is rather large and beyond the scope of this document, but most systems should include the ability to monitor and graph bandwidth. A sample bandwidth graph is shown in Figure 6.

Monitoring bandwidth allows a network engineer to spot trends and potential problems. For example, Figure 7 shows an internet connection that was using its maximum bandwidth of 250 Mbps for several hours; this appears in the graph as a distinct plateau. Users on this network might experience slowness or errors as they try to access the internet.

Figure 6: Daily Bandwidth Graph Showing Download (Blue) and Upload (Red) Utilization

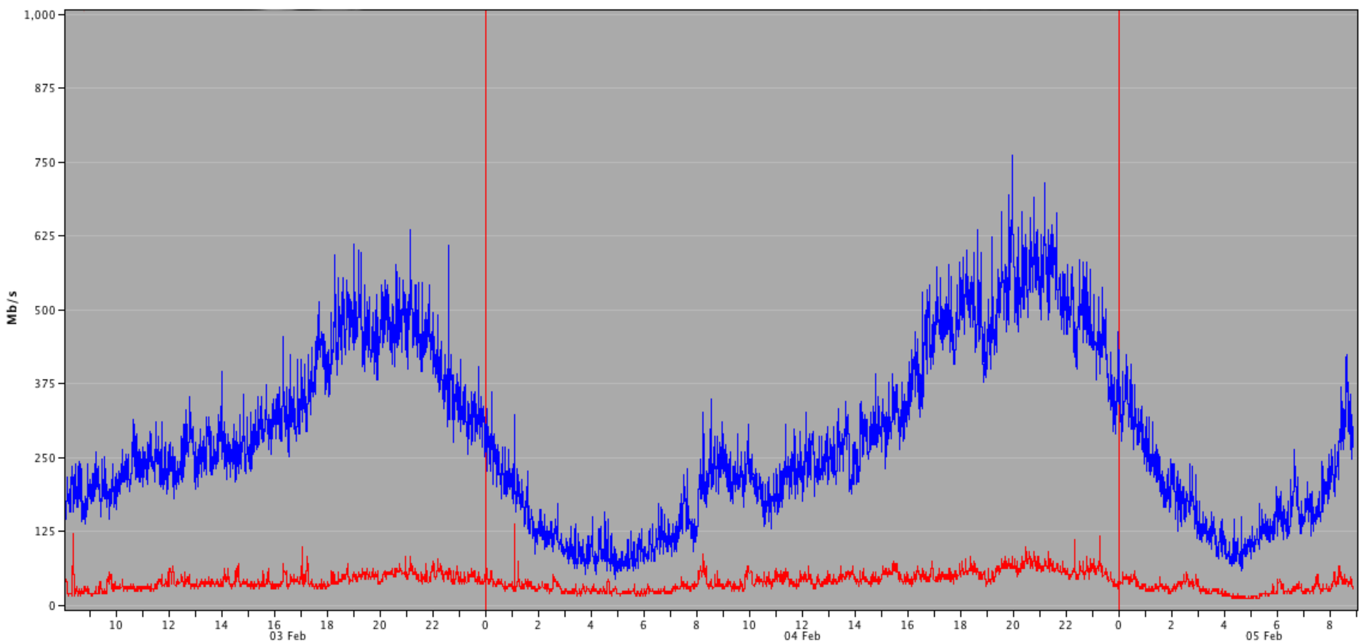

Figure 7: Bandwidth Graph Showing a "Maxed-Out" Connection

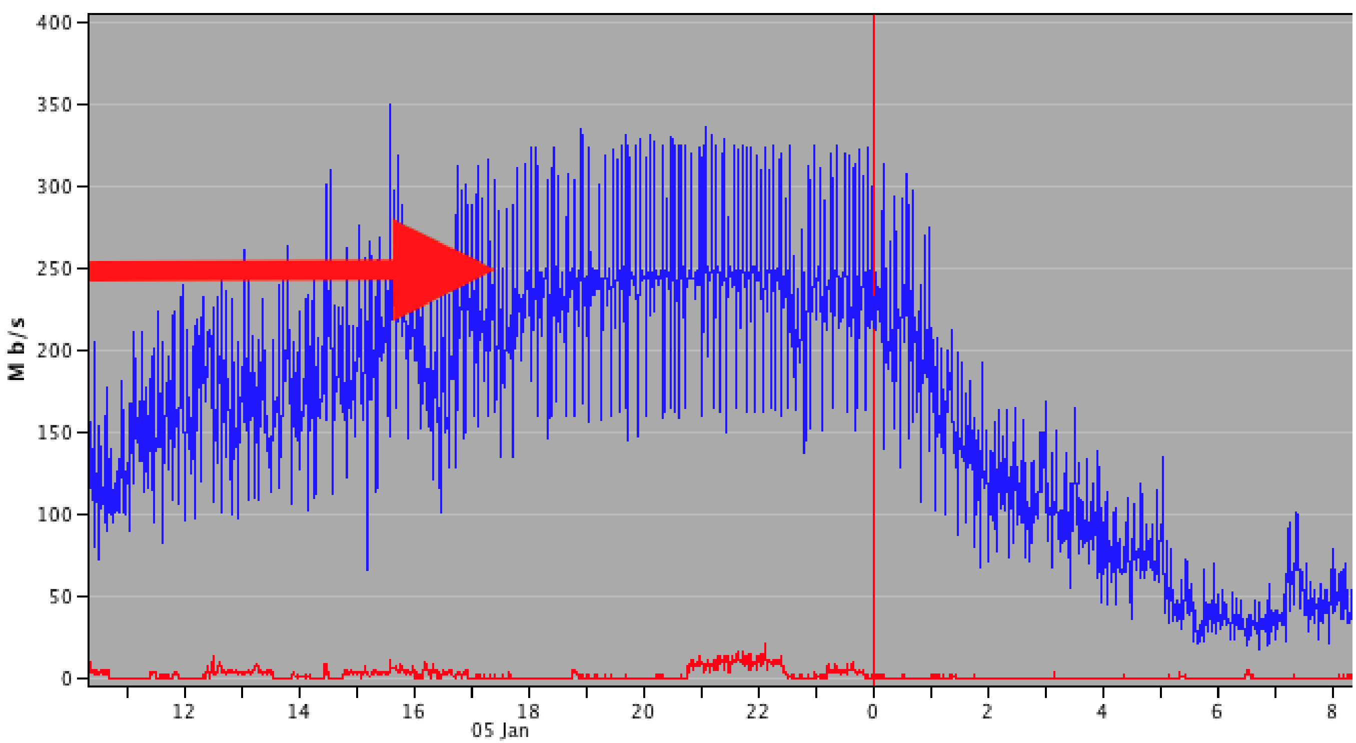

5 Collecting Usage Information to Determine Requirements

Depending on the purchasing process used, network specifications go into a bid document or other purchasing mechanism. With a well-formulated request for proposals (RFP), a library can solicit clear responses that can be evaluated and compared. The library can determine benefits and risks to each approach, whether a given response will meet current needs, and whether it will scale to meet future needs.

Below, we outline an approach that enables a library to derive its specifications based on high-level observations, viewed through network equipment and potentially a survey of users. This approach for bandwidth analysis and projection is also presented in the CTC document "The Library of the Future and Its Need for Bandwidth," and can be used by any library.

This approach is a "bottom-up" approach that estimates the specifications based on the characteristics of the library, the current usage, and the usage that is needed in both the short and long term.

We recommend starting with the following questions:

- How many users and devices are on the network?
- What is the type and amount of data usage?
- What is the flow of the data?

These questions lead to an estimate of the external network capacity that a library would obtain from a service provider or provision through other means, as illustrated in Figure 8.

Figure 8: Gathered Information Guides Specification of Needs

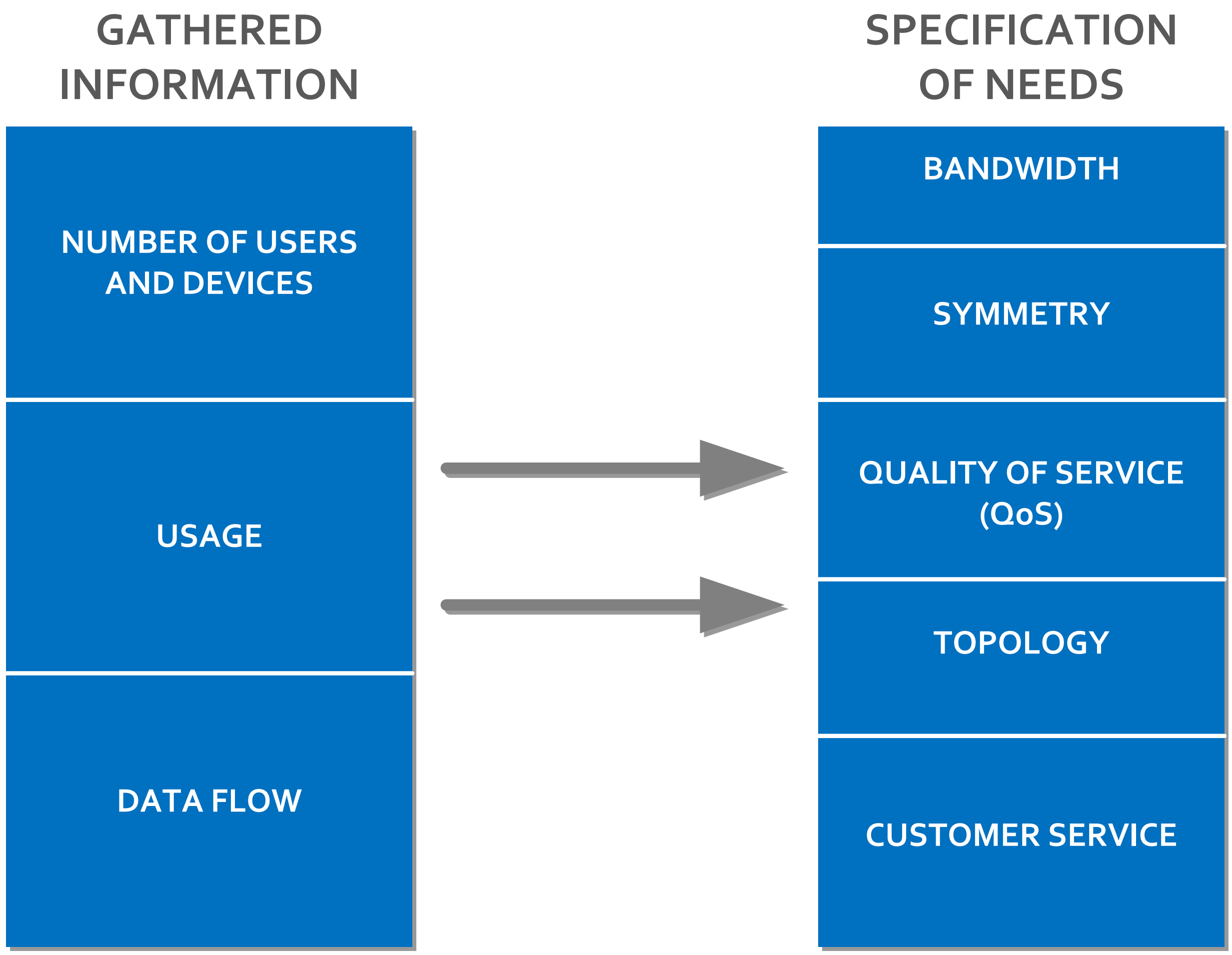

Other factors beyond the external network connection are also critical to meeting the library's broadband needs. These factors include building wiring, indoor wireless (Wi-Fi), power, and maintenance and operations of the network.

5.1 Number of Users and Devices

The first key number is the peak number of users on the network. One technique to obtain this metric is to take a manual count at a few peak times. It is important, however, to consider variations in the type of network usage and special events. If there is a meeting room, include the capacity of the room in this metric. It is also important to remember that the majority of patrons will be bringing their own devices.

Next, total the library-owned fixed devices and points to plug devices into the network. Count any smart TVs and streaming devices, conference/meeting equipment, administrative/internal devices, printers, scanners, CCTV, and building automation equipment (Figure 9).

Figure 9: Considerations When Developing a Total Number of Users and Devices

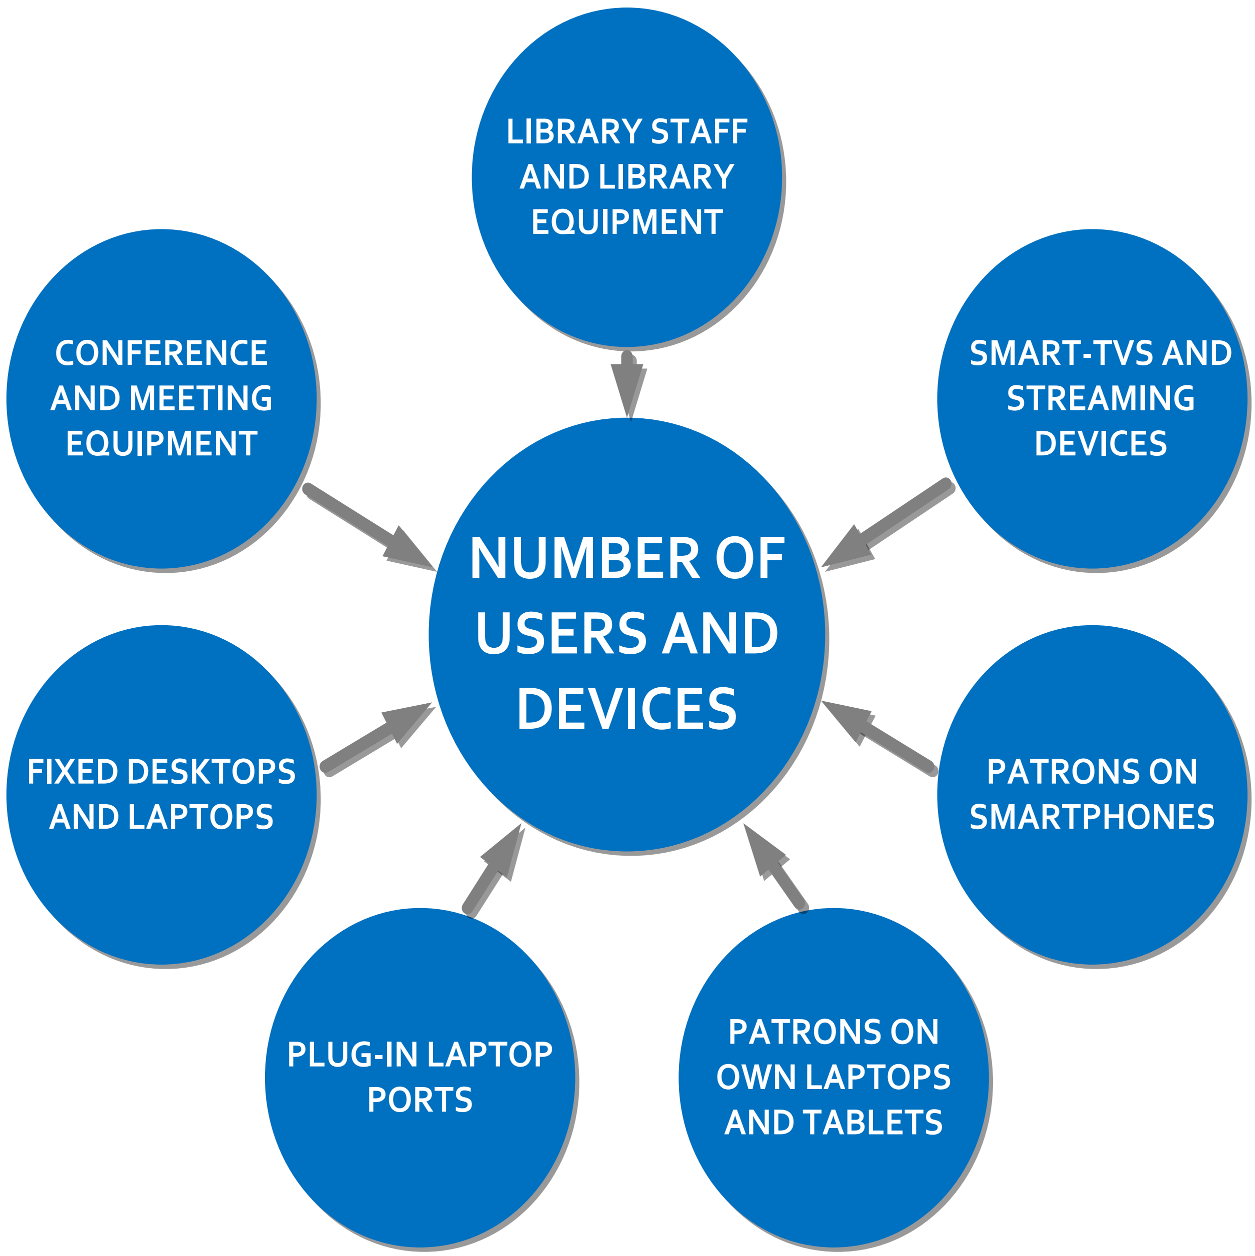

Finally, compile the information in a spreadsheet with a row for each category of users and devices, along with any assumptions. If there is a plan for expansion or other changes, consider creating columns for the current and planned states.

5.2 Network Usage

The next critical parameter is the degree and type of usage of the network. Measuring network usage can be a complex task, as usage changes moment to moment. Different types of patrons visit at different times of day with different needs and interests. Over a relatively short period of time, new devices and applications appear (e.g., tablets, Pokémon Go) that drastically change what is needed of the network.

As we discuss in Section 3, we encourage an approach that is more anticipatory of the applications that patrons need and less focused on the needs of the moment (which may change many times over the lifetime of the network and may themselves be a result of the users limiting what they want to do on an overloaded network). Therefore, we recommend against primarily basing specifications on the current usage of bandwidth. Rather, we recommend determining the usage based on the types of applications that patrons and the library use and anticipate needing to use. Example applications are shown in Figure 10.

Figure 10: Example Applications and Uses of Network Connectivity Services

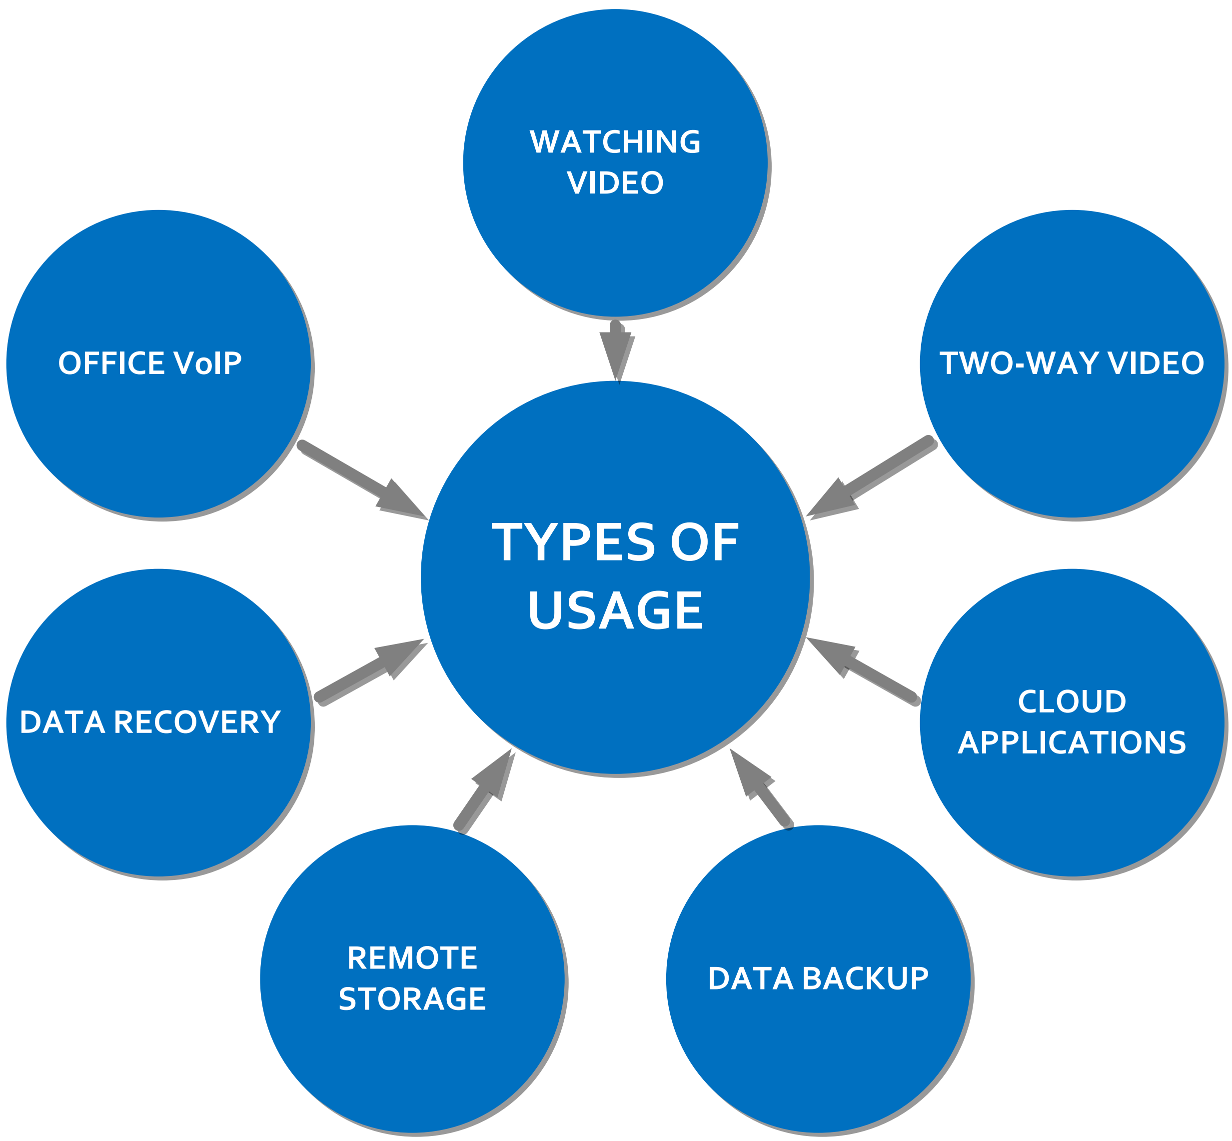

For an individual library system, we recommend determining these parameters through a range of strategies, including monitoring usage at a boundary firewall/router, conducting a walk-around survey at different times of day, reviewing strategic planning by the library, reviewing current and likely future applications used at the library, and conducting a survey of patrons and staff. This process will determine the approximate capacity needed per user, as well as the requirements for symmetry and QoS. It will also determine the degree to which other network devices and design considerations may be needed, such as locally storing (or "caching") frequently used content or media.

To document types of usage, add another tab to the spreadsheet created for users and devices. In this tab create a list of applications and uses. Note for each category if the usage appears frequent, moderate, or light.

Table 1 provides a range of applications, their typical capacity requirement, and whether the application communicates in "burst" or "continuous" communications.

Table 1: Applications and Typical Capacity Requirements

Application	Typical Bandwidth (Mbps)	Usage Pattern
E-mail	> 0.5	Burst
Internet browsing	0.5–1	Burst
VoIP	> 0.5	Continuous
Steaming audio	> 0.5	Continuous
HD video	3.5–5.0	Continuous
Videoconferencing	0.3–8.0	Continuous
Telepresence (participating)	0.5–4.5	Continuous
Telepresence (hosting)	3.0–16	Continuous
Virtual reality	50–500	Unknown
Download file in under 5 minutes		
50 MB file	1.5	Burst
100 MB file	2.5	Burst
512 MB file	12	Burst
1 GB file	23	Burst

5.3 Data Flow

Finally, the library needs to understand the flow of data on the network. Questions that must be answered include:

- Does all the traffic from branch libraries travel through the main library (or another single connection point)?

- Are there applications hosted in the central library for the system?
- Is there content filtering at a single point?
- Is there significant traffic to other institutions such as particular schools or universities?
- Is there a regional network aggregation point that is common to the libraries and all the important local institutions and the internet service providers?
- How symmetrical are these data flows?

Data flow can be observed by viewing utilization on network routers and firewalls. Data flow to specific branches can be determined by inspecting the packets at a router or firewall interface, though not all network equipment can easily perform this analysis. Although most data traffic in libraries is from the internet to the patron, it is important to understand how this may change to engineer the network for the future.

Figure 11 provides a high-level logical diagram of the type of locations that may be part of data flow to and from libraries.

Figure 11: Measuring Data Flow in a Library System

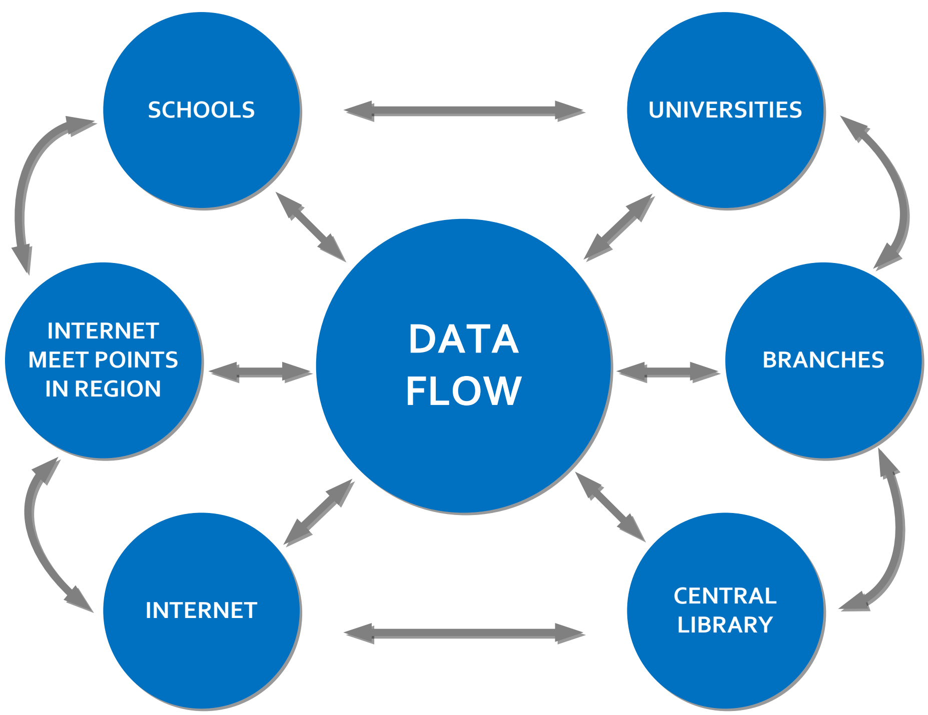

To document data flow, create a logical diagram that includes your main library and any branch libraries, the internet connection or connections, and any connections to other institutions that play a significant role on the network (such as a service provider, application provider, or partner).

If the libraries are on a regional network that includes educational or governmental institutions, include those institutions. If there is a regional network interconnection point that serves the libraries and those institutions, add that as well. If there is aggregation of network traffic at a specific location or filtering, add where that process take place. Include the network switches/routers and firewalls and their interfaces, and the speeds of those interfaces.

6 Developing Broadband Service Specifications Based on Requirements

Having determined the number of users and devices, the usage, and the data flow, a library is positioned to derive the specifications to fulfill its internet and network needs (see Figure 8, above).

Below, we provide an example of how to derive each set of specifications. This approach generally focuses on the most bandwidth-intensive and QoS-sensitive applications. Once those applications are properly supported, the other applications will be well served.

In a typical library, network use varies widely over the course of a week and increases over the course of a year. And the demands put on library networks are very likely to rise sharply given that societal data demands are growing. Cisco estimates that North American fixed internet traffic (this excludes traffic to smartphones and other mobile devices) will increase at a compound annual growth rate of 20 percent from 2016 to 2021; the worldwide estimate is even higher, at 23 percent.[9] In a multi-year contract, then, it is important to keep estimates high. Also, as discussed in Section 2, certain physical technologies (particularly copper) will severely limit total internet connectivity speed, and it is costly and time-consuming to install new infrastructure.

If planning for a four- or five-year contract, consider the networking world four or five years ago. The increasing adoption of devices such as tablets and smartphones, the availability of video streaming services such as Netflix and Amazon, and the development of increasingly sophisticated distance learning applications all came about within that period. Home internet speeds increased from a few Mbps to tens or hundreds of Mbps. The four or five years before that had comparable growth in broadband demand.

Therefore, it is reasonable to expect similar growth in broadband demand over the coming years, including growth to accommodate increased video resolution and usage of big data and cloud processing, to name a few foreseeable developments. Such trends imply that libraries may need multi-gigabit speeds in the coming years if they wish to provide state-of-the-art services to numerous patrons. While it is prudent to not overengineer and pay for services that are too far ahead of the state of the art, it is important to make reasonable preparations—striking the right balance.

[9] "Cisco Visual Networking Index," https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/complete-white-paper-c11-481360.html.

6.1 Bandwidth

Calculating how much broadband a library needs today requires determining two key variables: (1) the library's maximum number of concurrent broadband users, and (2) the amount of bandwidth each user needs.

Minimum Bandwidth Ceiling = [Number of Peak Users] x [Bandwidth Per User]

To estimate the number of peak users, consult the spreadsheet tab on Users and Devices. Total the devices. To that total, add about 50 percent of the peak number of patrons and staff in order to include usage of Wi-Fi equipped phones, tablets, and laptops.

If there are 40 desktops and laptops for staff and the public, as well as two videoconferencing units, the number of devices is 42. If the peak number of patrons and staff is 120, add half that—60—to the tally of peak users. **This brings you to a total number of 102 peak users.**

For bandwidth per user, consult the Usage tab to determine the dominant applications. Review the highest bandwidth applications on the list. If video is in the group, it is almost always the dominant application, as it is continuous; it dominates in environments where the other applications are bursty (browsing), slower (voice), or running only part of the day (data backup). Further, almost any device can become a video viewer or source, so the total bandwidth can quickly scale with video alone.

In a typical library environment, we recommend a ceiling based on enough bandwidth for users to take fuller advantage of what the internet has to offer today—what we call "adequate" bandwidth. Thus, in an environment where video is used by a substantial percentage of users, we believe a good benchmark for per-user bandwidth is 5 Mbps.

At 5 Mbps, HD video can be streamed, all online applications should work well, and a 1 GB file can be downloaded in less than half an hour. Netflix publishes a recommended speed of 5 Mbps per stream for HD-quality video.[10] Microsoft recommends between 2 Mbps and 8 Mbps for a Skype videoconference, depending on the number of people on the call.[11]

Not every device connected to the network will play high-definition video simultaneously, but some patrons will likely use more bandwidth than expected. A single device might be used to watch streaming video while also browsing the web and even downloading software updates in

[10] "Internet Connection Speed Recommendations," https://help.netflix.com/en/node/306 (accessed March 2017).
[11] "How much bandwidth does Skype need?" https://support.skype.com/en/faq/FA1417/how-much-bandwidth-does-skype-need (accessed March 2017).

the background. A single patron may use more than one device at the same time—for example, a tablet and a laptop—to participate in a videoconference call while downloading large files.

In addition to network usage by patrons, the library's bandwidth requirement must also account for increased usage by staff as resource catalogs, research, payment systems, and other internal applications continue to move online. If a library experiences a high-capacity day on the patron side, it is important to make sure library operations will not be impacted.

Using the equation above, we recommend a minimum bandwidth ceiling of:

102 x 5 Mbps = 510 Mbps

Some libraries have servers that are accessed by users offsite or have devices that "mirror" their content offsite, or even perform data processing tasks. The administrators of these networks typically have a good understanding of the total bandwidth requirements of such applications, and appropriate requirements for both the staff and users at the library should be added to the 510 Mbps above.

If users primarily browse and access little or no continuous media (e.g., streaming video), the per-user capacity can be lower, potentially in the 1 Mbps range. This is, to a large extent, because non-media applications are not continuous, so the network can support many simultaneous users on such applications.

In this environment, the minimum bandwidth ceiling would instead be

102 x 1 Mbps = 102 Mbps

We note, however, that this environment will become starved at peak times—resulting in poor network performance—if any appreciable number of patrons start using video or media applications. If this level of bandwidth were chosen, there would need to be a way to increase bandwidth as user needs increase.

While we believe a bandwidth of 5 Mbps per user would serve most environments for the near term, over the next few years we expect that number may double, based on past trends.

To accommodate the above scenarios, we would recommend the library specify the following:

500 Mbps capacity with the ability to scale to 1 Gbps without additional construction

We would also recommend that respondents to an RFP be required to provide the cost for providing service at higher speeds, such as 100 Mbps increments up to 1 Gbps, as well as any cost for changing the service.

Sometimes a service provider will propose a burstable connection, for which the library pays a base rate. This ability to "burst" to higher speeds as needed may be a suitable solution as long as any incremental burst cost is reasonable and the mechanism for using the burst (automatic as needed, or call-ahead if high demand is anticipated) is understood and can be effectively managed by the library.

The library should include in its estimated needs any new services that it anticipates deploying over the duration of the contract. The American Library Association and the FCC have adopted a target goal for libraries based on the number of people they serve. Libraries that serve fewer than 50,000 people should have at least 100 Mbps. Libraries that serve more than 50,000 people should have a 1 Gbps connection.[12]

6.2 Symmetry

The symmetry needed for a network connection depends on the applications used on the network, as well as whether the library hosts content, uses cloud applications, or performs intensive data backups. Typical browsing and web content viewing has a downstream to upstream bandwidth usage ratio of 10:1. In a business environment, a ratio of 5:1 is common to support remote hosting, emails, and videoconferencing. In an enterprise setting with cloud applications, a substantial amount of videoconferencing, and real-time data backup, the network use becomes almost symmetrical.

It should be noted that based on the applications above, most institutions are in fact becoming more symmetrical in their network use over time. For all but the smallest libraries, a ratio of 5:1 is a recommended starting point, with the ability to scale to a fully symmetrical network.

For the above scenario, then, the recommended specification is:

500 Mbps downstream, 100 Mbps upstream, with the capability to scale to 1 Gbps upstream and 1 Gbps downstream without additional construction

We also note that increases in symmetry potentially require increased cost, particularly in small library environments. This is because these libraries often have cable modem and DSL technologies that are suitable for a highly asymmetrical environment but cannot support a symmetrical environment.

---

[12] FCC, "Summary of the E-Rate Modernization Order," (https://www.fcc.gov/general/summary-e-rate-modernization-order), Order pp. 34-38. See also "Comments of the American Library Association," GN Docket No. 16-245, Sept. 2015, https://ecfsapi.fcc.gov/file/1090767678599/American%20Library%20Association%20Section%20706%20Comments%209_7_16.pdf.

If specifying services for a small library, where the long-term need may not exceed the low 100 Mbps range, it is worth requesting pricing for a range of symmetry levels and considering the cost tradeoffs.

6.3 Quality of Service (QoS)

We recommend that libraries specify QoS parameters that suit the applications used and that set an expectation for the availability of their networks—assuming that the increased cost for QoS is manageable. Typically, only business- and institution-grade connections have QoS requirements and penalties for noncompliance.

Typical availability for a large critical institution will be 99.999 percent (exclusive of fiber breaks); for a smaller institution it would be 99 percent (exclusive of fiber breaks). QoS terms also include a time to restore (excluding fiber), which is in the few-hours range, and may be 24/7 or limited to business hours.

Fiber breaks are typically excluded and handled separately—for example, with a committed time to respond on site (e.g., four hours)—but repair time is dependent on the nature of the problem and the need to coordinate with traffic and the other utilities. For a highly critical facility, a library may specify two diverse routes or may contract with two separate service providers for the route.

Parameters for other QoS terms that are typical in a high-speed institutional connection, and that would be suitable for the above scenario, are:

Latency (one-way) <5 ms

Jitter (one-way) <1 ms

Packet loss (one-way) < 0.3%

These are adequate for most internet applications, including real-time carrier-grade voice and video.

Penalties are usually on a sliding scale, with progressively higher levels depending on the length and severity of an outage, generally reaching up to a full month's credit. For repeated failures, the library may have the ability to terminate the agreement.

The library may specify the parameters in an RFP. Some institutions, especially in smaller metropolitan areas or rural areas, do not specify QoS in requests for a bid but require the provider to respond with the terms they can provide, enabling the library to compare among respondents or negotiate.

6.4 Topology

The topology in a specification is generally fixed by the structure of the library system and affiliated institutions. Connecting libraries in a system to the central library and centralizing the internet connection allows a library system to consolidate functions such as filtering and security, creates economies of scale, and allows for more flexibility in operating internal network services.

It is also a good idea to look at opportunities to aggregate network connections and engage in partnerships with other institutions. For example, if the libraries can connect to a regional fiber network operated by the local government, an educational consortium, or the state, it potentially can aggregate demand among all the libraries, plus the other institutions.

This may replace a smaller network purchase with a significantly larger one, creating economies of scale. Large-scale internet service is not only cheaper; it also draws the attention of larger, higher-tier service providers that potentially offer higher-quality connections and are closer to the internet backbone. If the regional network connects to a local or regional internet access point, there is also the potential to directly connect with other service providers or institutions such as research networks and universities present at the location. This can reduce the cost and increase the quality of the connection.

6.5 Defining Network Requirements

With an understanding of the library's current and future bandwidth needs, the staff can develop broadband service requirements. This step is especially important if the library intends to use a competitive bidding process or issue a request for proposals. If the library intends to apply for E-rate funding, all proposals must go through a formal evaluation process that uses cost as the most important, though not the only, factor. Even outside the E-rate process, it is important that the library set requirements that will ensure its needs will be met but are flexible enough that good solutions will not be disqualified.

One way to approach this is to focus on functional requirements of the network rather than strict technical requirements. Requirements such as speed, uptime, delay, jitter, and loss, can be treated as minimum requirements below which a proposed service would be disqualified.

Table 2 shows requirements for an example library system with three locations. The requirements specify location, minimum speed, critical applications, and the number of computer workstations at each library. Other requirements such as acceptable delay, jitter, and loss for videoconferencing; uptime; resiliency; and average bandwidth usage for critical applications may be spelled out elsewhere in the requirements document.

Table 2: Example Bandwidth Requirements

Facility	Location	Minimum Connection Speed	Purpose	# Computers
Main Library	123 Main St.	1 Gbps downstream/ 1 Gbps upstream	Library management software, email, videoconferencing, streaming video, internet access	75 public, 20 staff
Branch Library A	200 Oak St.	100 Mbps downstream/ 100 Mbps upstream	Library management software, email, streaming video, internet access	15 public, 5 staff
Branch Library B	300 Elm St.	100 Mbps downstream/ 100 Mbps upstream	Library management software, email, streaming video, internet access	15 public, 5 staff

These requirements allow potential respondents to propose a variety of technical solutions, thereby giving the library more options, while explicitly stating what the library considers an acceptable level of service. Respondents who likely have differing pre-existing network infrastructures also have the freedom to propose the most cost-effective network topology based on construction costs and other factors that the library would not be aware of. A respondent may propose, for example, connecting the three library locations with a WAN and a shared connection to the internet or, alternatively, a completely separate internet connection to each location. In short, the library evaluates its current state and where it needs to go while the service provider proposes how to get it there.

6.6 Defining WAN Requirements

If the library's solution includes a WAN, the staff can use the existing network requirements and data flow information to develop requirements for the connections between individual locations. In general, any requirements placed on the library's internet connection (such as acceptable delay, jitter, and loss for videoconferencing; uptime; and resiliency) should be met or exceeded by all WAN connections. Significantly higher performance or a larger amount of bandwidth may be required, depending on the design and intended use of the WAN.

A WAN link will need enough bandwidth to accommodate the data that flows between the two connected locations as well as the internet traffic to any location that relies on that link to reach the internet (Figure 12).

Figure 12: Bandwidth Requirement for WAN Links

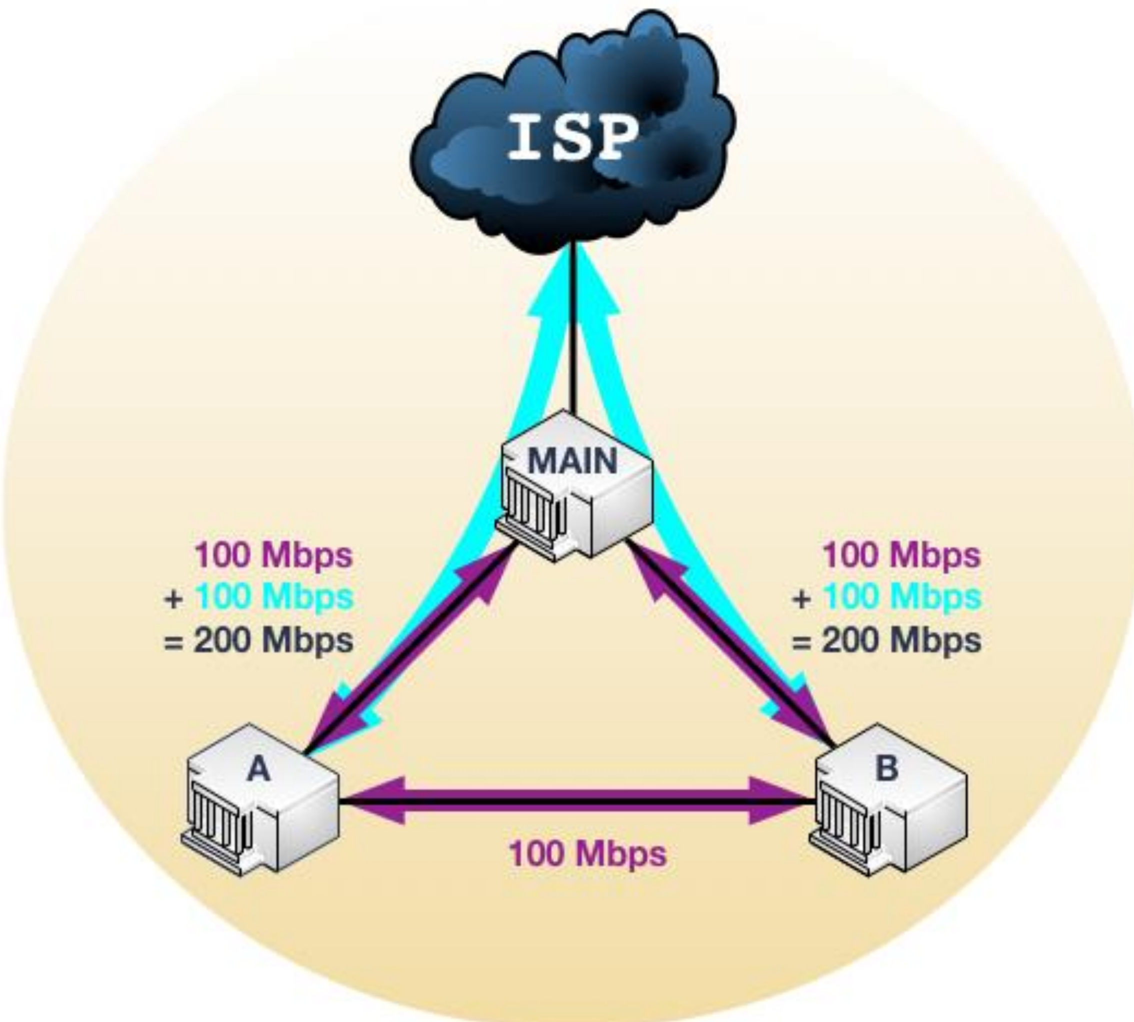

Further, if the WAN is designed to provide redundancy in the event that one of the links fails, each link will need to accommodate the additional data that will flow over it during a failure (Figure 13).

Figure 13: Bandwidth Requirement for WAN Links in the Event of an Outage

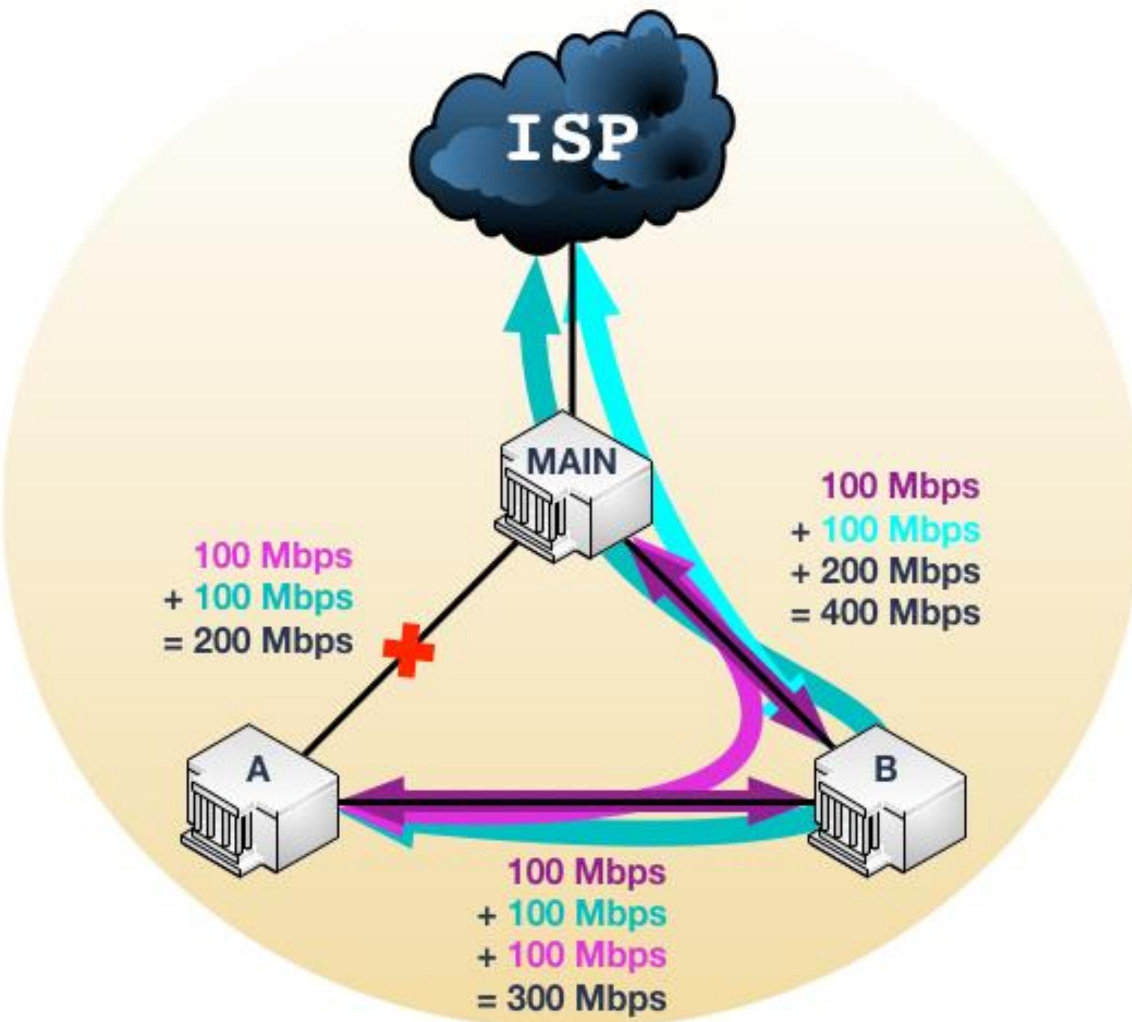

In addition to internal data flows and redundancy, it is important to allow room for growth when developing requirements for a WAN. Because a faster connection to the internet might change the way that staff and patrons use the network, a new or faster WAN may also change usage patterns and present opportunities for new services. The FCC does not set a specific goal for library WAN speeds, but the target bandwidth for school systems is 10 Gbps per 1,000 students, or 10 times the target for internet access.[13]

[13] "Summary of the E-Rate Modernization Order," https://www.fcc.gov/general/summary-e-rate-modernization-order (accessed April 2017).

The WAN design should also take into account the criticality of each location and the amount of bandwidth needed when specifying the topology of the WAN. Locations with higher bandwidth requirements or where connectivity is highly critical should be the first candidates for redundant connections and should have fewer links to cross in order to reach the internet. This will reduce the amount of traffic on the WAN between locations and reduce the likelihood that a network failure will result in the loss of services at other locations or to customers accessing library services offsite.

At the same time, the library could save money by connecting a less critical location to the WAN by a lower-bandwidth link to the nearest branch location, instead of a longer connection to the main hub. In Figure 14, for example, Branch C is less protected from link failure and could cause a traffic bottleneck at Branch A, but it would also be less expensive to connect.

Figure 14: Sample WAN Scenario

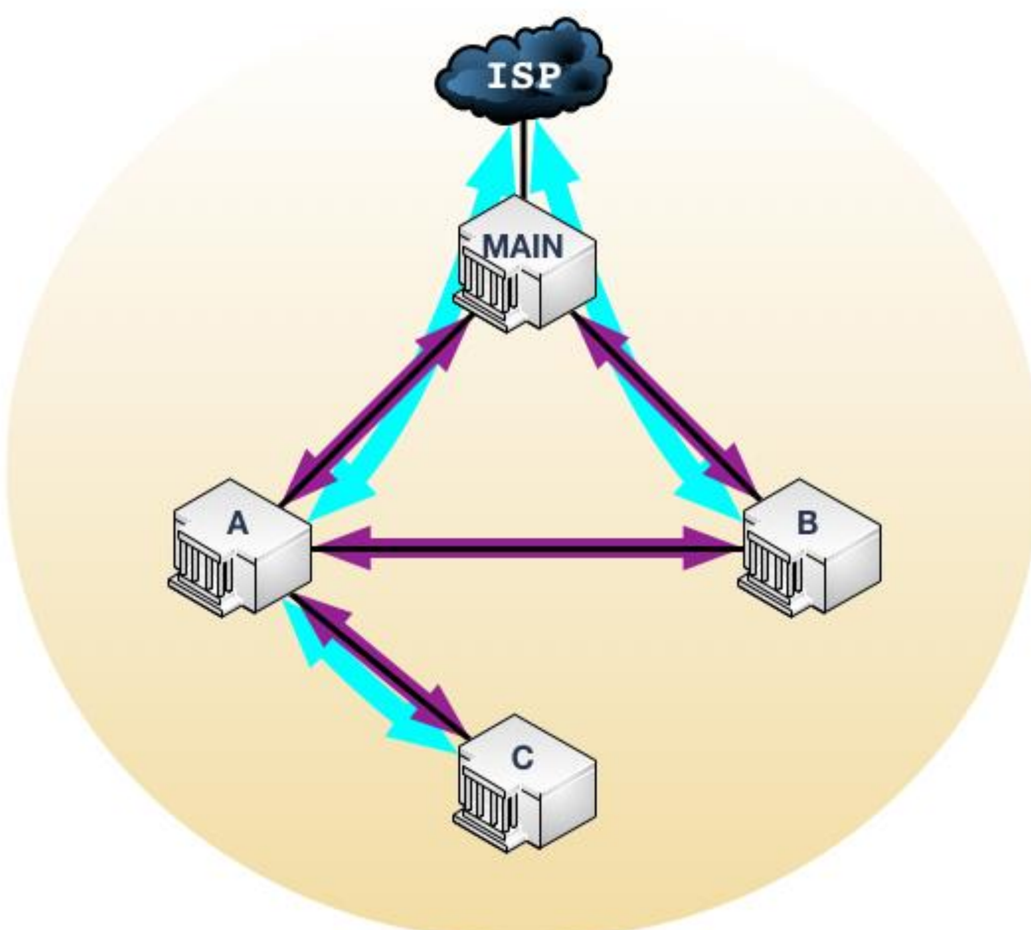

6.7 Project Management and Customer Service

In addition to requirements for the network itself, a successful RFP should include requirements for project management services and customer service that the library will receive from the internet service provider. RFP responses should list key personnel who will be involved in implementation of the service, including any required construction, installation, and testing.

The library should request an implementation plan for the project, including detailed steps and milestones, relationships, and dependencies among tasks, how progress will be tracked and reported, and a narrative describing how the project will be completed within the desired time frame. If there are approvals or processes that the provider must consider, those should be outlined in the RFP so the respondent can take these into consideration.

The RFP should also require a description of the customer service the library will receive from the service provider. This should include response times for requests made by telephone, email, or the provider's website as well as response time for an on-site technician to arrive and the mean time to repair (MTTR) for minor and major issues. In addition, the service provider should provide procedures for contacting customer support and escalating issues as well as listing any individuals who will be assigned to the library as a point of contact. The library should specify any requirements for regular meetings or consultation.

6.8 Information to Request from Providers

When the library selects a response to its RFP, it is not just purchasing a product but entering into a contractual relationship that may last several years, and potentially building infrastructure that will last decades. In addition to requesting information about the service itself, it is critical to request information from the respondent that may factor into the library's decision.

The RFP should request a helpful amount of information without being overly burdensome so as to discourage potential bidders. Respondents should be able to show their compliance or ability to comply with any applicable rules for contracting and procurement as well as local permit and policy requirements. If the library requires compliance with any funding programs, such as E-rate, or wishes to give priority to certain types of respondents, such as a minority- or woman-owned business (MWBE), this must be clearly stated. The following are examples of questions the library may want to ask potential respondents:

- Full legal company name
- SPIN ID (or status of Form 498 application) if seeking E-rate funding
- Year business was established
- Number of people currently employed

- Most recent annual report, if a public company
- Mission and vision statements
- A description of the qualifications, experience, capability, and/or capacity of the company to successfully provide the broadband service and complete the project in a timely manner
- A description of the qualifications of the members of the proposed project team that will be assigned to the project
- Location and description of current network operations
- Information on current broadband clients including:
    - Total number of current clients
    - A list of broadband services provided to similar clients
    - Evidence of successful completion of a project of a similar size and complexity
- References: Contact information for three references from projects similar in size, application, and scope; and a brief description of their broadband installations
- Evidence of company's financial bonding status

Appendix A: Exploring Potential for E-Rate Funding

The federal E-rate program (also known as the Schools and Libraries Program) provides substantial discounts to schools and libraries for eligible telecommunications services and required components such as network equipment or construction. These discounts can enable libraries to provide beneficial services to their patrons and communities that would be otherwise cost-prohibitive and to keep pace with the growing demand for these services.

E-Rate Background

The E-rate program is a component of the Universal Services Fund (USF), which was established by the Federal Communications Commission (FCC). The fund is intended to promote the "principle that all Americans should have access to a baseline level of telecommunications service and further the public interest of keeping all Americans connected."[14] The Universal Service Administrative Company (USAC) is an independent, not-for-profit company established to administer the fund and ensure compliance with the program.

The program, as outlined in this document and in greater details on USAC's website, involves a competitive bidding process in which the library must select the provider with the most cost-effective proposal to meet its need.

Once the needs of the library system are understood, the process of applying for E-rate discounts and creating and publishing a request for proposals (RFP) can begin. Creating an effective RFP is an involved process, and the language and directions must enable the library to select a respondent that truly meets its current and future needs.

E-Rate Application Process

USAC requires that RFP respondents be evaluated fairly and that the most cost-effective solution is chosen. The RFP should be flexible enough to allow for creative responses yet have clear requirements. Speed and cost are only part of the equation. Factors such as service reliability, construction, permitting, and staffing needs all play a part in determining the success and feasibility of the selected service.

The RFP process must be timed carefully with the E-rate application process, which involves filing multiple FCC forms with USAC, which will carefully review your request to ensure it meets requirements for eligibility and fairness. In our experience, sufficient time is one of the most critical factors to completing this process successfully. Applicants must allow themselves time to complete and file the required FCC forms, to craft the RFP and an RFP evaluation, to correspond

[14] "Universal Service," http://www.usac.org/about/about/universal-service/default.aspx.

and negotiate with RFP respondents, to correspond with USAC during the application review process, and to work with their own procurement process to obtain the selected services.

USAC generally opens the E-rate discount application window approximately six months prior to the start of the funding year and keeps the window open for about two and a half months[15] (funding years begin July 1 and end June 30 the following year). Competitive bidding on an RFP must remain open for at least 28 days and a discount application for funding is submitted, meaning that RFPs must be submitted at least 28 days before the close of the application window.

USAC generally begins accepting the form (FCC Form 470) required to open the competitive bidding process 12 months before the beginning of the funding year, and RFPs can be submitted as soon as that form is made available. This gives applicants over seven months to submit an RFP, correspond and negotiate with service providers, evaluate and select an RFP response, and apply for E-rate discounts.

Before beginning the RFP process, the library should also take time to understand the E-rate process and forms, decide on governance and staffing, gather required information, assess current and future bandwidth needs, and research potential solutions. The library may also enlist the help of higher-level agencies that have completed the E-rate process before or of a consultant who can help the library through the process. This outside help is often essential to completing the E-rate and competitive bidding process in an effective and timely manner.

USAC provides detailed information about the application process on its website at www.usac.org/sl/applicants. We provide a brief overview here.

1. Preparation
2. Submit Form 470 with RFP
3. Bidding (28-day minimum)
4. Select a provider and purchase services
5. Submit Form 471 during application window
6. Start receiving service (July 1)
7. Submit Form 486
8. Invoice USAC

Preparation

Time may be the most critical factor in a successful E-rate application. The library must allow itself time to construct an effective RFP, communicate and negotiate with respondents, evaluate proposals, select a provider, and execute a contract with the chosen provider in time to submit

[15] "Step 3: Applying for Discounts," http://www.usac.org/sl/applicants/step03/default.aspx.

Form 471 during the application window. The most critical time is between the release of Form 470, which generally occurs approximately one year before the funding year, and the closing of the Form 471 application window, which generally opens approximately six months before the funding year and closes approximately two and a half months later. This allows a maximum of eight and a half months to complete the RFP process and apply for E-rate discounts.

A library can take several steps to prepare for the RFP process before Form 470 is released, to ensure that it has the necessarily time and resources for the actual bidding and selection process. We recommend that the library decide on governance and staffing for the E-rate process early on. This will include deciding whether the library will apply on its own or as part of a consortium with other libraries and schools, as well as appointing the staff or committees that with be primarily responsible for the E-rate application.

The appropriate staff should review the eligibility requirements for the library and ensure that each participating entity has a billed entity number (BEN). They should also ensure that the library or consortium has an E-Rate Productivity Center (EPC) account, which will be the center of activity and communication during the E-rate application process. Note that the library must comply with the Children's Internet Protection Act (CIPA) in order to be eligible.

To budget enough time for negotiation and execution of the contract, review any internal budgeting, purchasing, or approval processes that may be required. Make sure to allow ample time between the close of the bidding process and the deadline for the Form 471 application for your organization to execute the contract with the chosen provider and order the agreed-upon services.

This would also be a good time to evaluate the library's needs and decide what it will ask for in the RFP as well as the method it will use to evaluate responses. It is important to review the Form 470 (which is updated annually by USAC) and any changes in eligible services announced for the new funding year. But much of the groundwork can be completed before this information is released, even if the library does not start writing the RFP.

Form 470 and the Request for Proposals

Filing Form 470 along with your RFP will officially open the competitive bidding process. Bidding must remain open for at least 28 days, which means that Form 470 must be submitted at least 28 days before the close of the funding application window. However, we advise allowing much more time between closing the bidding and submitting the funding application, so Form 470 should be submitted much earlier.

The RFP should be a clearly written description of the library's desires and requirements, including requirements for maintenance, any required construction, and metrics for acceptable performance in addition to the service itself. All proposals must be evaluated fairly and must use cost as the primary, though not only, factor. Because of this, it is important to write the RFP in such a way that successful responses will accomplish the library's goals and that desirable proposals will not be unnecessarily disqualified.

Note that inexpensive internet access services that meet certain requirements are exempt from filing Form 470 and going to bid. Libraries should review the requirements for the current funding year and explore available internet access options to see if any that would meet the library's goals and needs would qualify.

Competitive Bidding

As stated previously, competitive bidding must remain open for a minimum of 28 days. However, the library will want to plan for a bidding window that allows enough time to circulate the RFP and for bidders to ask questions and craft their responses. The amount of time required will vary depending on the amount of information required by the RFP.

Selecting a Provider

Once bidding has closed, the library must select a provider by fairly evaluating all responses and selecting the most cost-effective solution. USAC provides guidance for conducting a formal evaluation. We recommend defining the evaluation process within the RFP itself to ensure the RFP clearly reflects the items by which responses will be graded and that appropriate information is provided to all respondents.

The library should allow time to discuss the details of the service with the provider, to review the contract if there is one, and to complete any required internal purchasing and approval procedures. The library and the provider should also discuss who will be responsible for billing USAC for the amount of the discount. Note that services should not begin before the funding year, but some components required by the service may be eligible for installation before July 1. The library will need to review the specific requirements for the current funding year.

Apply for Discounts Using Form 471

Form 471 details the services the library will receive and the discounts it is eligible for. The window for filing this application generally opens approximately six months before the funding year and remains open for about two and a half months. The form requires several pieces of information about the library as well as the chosen provider's Service Provider Identification Number (SPIN).

It would be beneficial to review the requirements for this form during the preparation stage to allow adequate time for the information to be gathered. It would also be beneficial to inform potential RFP respondents that they will be required to obtain a SPIN if they do not already have one.

The library's applications are reviewed in the Program Integrity Assurance (PIA) process to ensure all applications are program-compliant. The library must be prepared to respond to any questions from the reviewers within 15 days. Once the PIA process is complete, the library will receive a Funding Commitment Decision Letter (FCDL).[16]

Confirm Service and CIPA compliance with Form 486

Within 120 days of either receiving the FCDL or starting services (whichever is later), the library must submit Form 486 to confirm that it has received service and is complying with CIPA.

Invoice USAC

After Form 486 has been approved, either the library or the service provider may invoice USAC for the amount of the discount applied to the services.

[16] We note that applying for funding from the E-rate program comes with risk. There are no guarantees that funding will be awarded.

Appendix B: Technologies and Services Associated with Lit Fiber

Dense Wave Division Multiplexing (DWDM) – DWDM allows for many distinct connections to be transmitted over the same fiber, as different wavelengths (colors). DWDM can drastically extend the capacity of a single strand of fiber and allows for completely separate channels of communication with each customer or service being assigned its own "wave."

Transport services – Service providers may advertise the type of transport service used on their networks. Three transport service technologies are in use.[17] Regardless of the technology used for transport, the library will most likely use a standard Ethernet connection to connect to the service provider's equipment. The type of transport network may affect available features.

Dedicated Internet Access (DIA) – DIA provides connectivity from the internet to a library via a variety of kinds of point-to-point connections that are not shared with other customers. These include:

1. ***Ethernet Private Line (EPL or EVPL) –*** An Ethernet Private Line provides a point-to-point service connecting two locations. An EPL provides a secure private connection suited for high bandwidth, low latency, or sensitive data applications. An Ethernet Virtual Private Line (EVPL) uses additional software to allow for multiple services between multiple sites. (This additional functionality is known as a virtual local access network or VLAN.)

2. ***MPLS Layer 3 VPN (MPLS IP VPN) –*** An MPLS Layer 3 VPN uses the service provider's MPLS network to connect multiple sites together. The service provider handles routing between the sites. Because these services are scalable and can support very high-quality levels of service, they are popular for interconnecting branch locations or connecting branches back to a central data center. They are also one of the most widely deployed technologies for business WANs. Service providers sometimes advertise these services as simply "MPLS."

3. ***Virtual Private LAN Service (VPLS) –*** VPLS acts as an extended LAN or Layer 2 network where the library would be responsible for routing between the connected sites. VPLS is comparable to MPLS in that both provide a private network with QoS capabilities, but VPLS does not scale well compared with MPLS when connecting a large number of sites.

[17] Multi-Protocol Label Switching (MPLS) and Synchronous Digital Hierarchy/Synchronous Optical Networking (SDH/SONET) are the most commonly used technologies, though Carrier-Ethernet Transport (CET) is growing in popularity. SDH/SONET evolved from telephone network protocols, while MPLS and CET are both based on modern packet-switched networks.

IPsec VPN – IPsec VPNs use the IPsec protocol to establish encrypted communications between sites or to a central point on the service provider's network. IPsec VPN traffic is routed along with other internet traffic and may not offer the same performance and reliability as other WAN technologies, especially if not all library sites connect to the internet through the same service provider. An IPsec VPN does provide flexibility to be used among multiple locations regardless of the type of connection or the service provider, if the tradeoffs in management, performance, and reliability are acceptable.

Ethernet (Carrier Ethernet or Metro Ethernet) – Ethernet has traditionally been used in LANs but has been extended for use in large carrier networks and private WANs. Using the same Layer 2 technology in the LAN and the WAN can simplify network operations and management. Ethernet service is comparable to VPLS in that the library would control routing over the WAN and that Ethernet is less scalable than MPLS for large numbers of sites.

Appendix C: Understanding the Differences Between Local and Wide Area Networks

The library's connection to the internet is critical, but other parts of the network will also impact the perceived performance of the network. Other potential causes of network slowness include local area network (LAN) ports, router/switch capacity, wireless access point (WAP) speeds, and Wi-Fi utilization (bandwidth utilization and number of connected devices).

Network equipment and cable that is old, malfunctioning, or simply insufficient can prevent patrons and staff from utilizing the full potential of a fast broadband connection. A wireless access point, for example, can act as a bottleneck between a patron's laptop and a fast internet connection. If the access point is using an older, slower Wi-Fi standard, suffers from radio interference, or simply has too many devices connected to it, the devices connected to it may have a slow and unreliable internet connection.

To ensure that library patrons and staff will benefit from a high-speed broadband connection, it is necessary to evaluate the parts of the network between the connection to the internet and the devices that people use to connect. Like connections to the internet, LAN and wide area network (WAN) technologies can qualify for E-rate funding.

The Local Area Network

The LAN is the network within the library building. For most libraries, the LAN is comprised of the routers, switches, wireless access points, computers, servers, and cables within the library building. The LAN is usually managed by library staff or a contractor as opposed to the ISP. The design of a LAN can vary greatly and is beyond the scope of this document. However, there are criteria the library can use to determine if the LAN needs improvement. These criteria take into account basic internet connectivity and may not meet the needs of every library.

Cable: Copper Ethernet cable should be rated Category 5 or higher, which will allow for gigabit speeds within the library. New installations should use Category 6 cable, which allows speeds up to 10 Gbps.

Device connections: Individual workstations with wired network connections should be connected to a switch at 100 Mbps or higher. New installations should use 1 Gbps connections.

Uplink ports: An uplink is the port that connects a switch to the rest of the network. For example, if you have five workstations connected to a switch, and that switch has one connection to the library's main router, the router connection would be the switch's uplink. This port should be no slower than the ports that connect to individual workstations and may need to be faster, depending on the number of connected devices and the amount of data they generally send and receive.

Wireless: There are many different Wi-Fi standards in use and different kinds of wireless access points have very different capabilities. Most current devices use the n or ac Wi-Fi standards. Slower standards like a, b, or g may not provide acceptable performance. Wireless performance can also be affected by coverage area, interference, the number of connected devices, and the speed of the connection between the access points and the switch.

A simple hardware inventory may reveal areas where the library's LAN could be improved, but it may also be helpful to monitor connections or perform tests at different points in the network. Tests and monitoring may reveal points where the network does not perform as expected, where traffic bottlenecks can occur at busy times, or where wireless coverage is not adequate.

The Wide Area Network

If the library is part of a system of libraries or other institutions, it may connect to other locations over a WAN. A WAN connects two or more geographically diverse locations and usually involves connections that have been leased from an ISP, unless the system has constructed its own fiber. It allows these locations to behave as if they are on the same local network, passing data between locations without going over the internet. A WAN by itself does not provide access to the internet, but it can allow multiple locations to share one or more internet connections and provide resiliency in case of failure.

Requirements for WAN connections vary depending on application, but they must account for data passing between library locations as well as any internet traffic in the case of shared internet connections. If, for example, two branch locations connect to the internet via a WAN connection to the primary branch or office, and are connected to each other directly, each connection to the primary office should be able to accommodate traffic from both branches at once in case one branch loses its connection to the primary office and must connect through the second branch.

The FCC has not set a target speed for libraries that are connected by a WAN, but it is tracking the number of libraries that have a WAN capable of providing 10 Gbps service.

Appendix D: Broadband Terms and Definitions

The following glossary defines commonly used terms as a quick reference for personnel evaluating different broadband solutions.

Access Fiber – The fiber in an FTTP network that goes from the FDCs to the optical taps that are located outside homes and businesses in the rights-of-way.

Broadband – High-speed internet access that is always on and is faster than dial-up access. In 2015, the FCC updated the legal definition of broadband to refer to services providing at least 25 Mbps download and 3 Mbps upload.

Burstable – An internet connection that can temporarily exceed the bandwidth speed that a subscriber is purchasing.

CPE – Customer premises equipment; the electronic equipment installed at a subscriber's home or business.

Dark fiber – Fiber that is unused or not connected to electronics that "light" the fiber. The phrase can also refer to fiber that is provided in a "dark" state to be connected and lit by the customer.

Dark fiber lease – A contract to lease dark fiber, typically for a shorter term than that in an IRU agreement, paid on a month-to-month or annual basis.

Drop – The fiber connection from an optical tap in the right-of-way to the customer premises.

FDC – Fiber distribution cabinet; houses the fiber connections between the distribution fiber and the access fiber. FDCs, which can also house network electronics and optical splitters, can sit on a curb, be mounted on a pole, or reside in a building.

FTTP – Fiber-to-the-premises; a network architecture in which fiber optics are used to provide broadband services all the way to each subscriber's premises.

IP – Internet protocol; the method by which computers share data on the internet.

IRU – Indefeasible right of use; an agreement, typically covering 10 to 20 years, under which the customer has the right to use dark fiber strands on a network.

ISP – Internet service provider; an organization that provides services enabling customers to connect to the internet.

OSP – Outside plant; the physical portion of a network (also called Layer 1) that is constructed on utility poles (aerial) or in conduits (underground).

QoS – Quality of service; a network's performance as measured on a number of attributes.

RFP – Request for proposal; a document that solicits price/solution proposals for a given project.

ROW – Right-of-way; land reserved for the public good such as utility construction. ROW typically abuts public roadways.

WAP – Wireless access point; a device that broadcasts an internet signal wirelessly to enable devices to connect to the internet.

Wi-Fi – Wireless fidelity; the networking technology by which computers and other devices transmit data wirelessly.