Chapter 11 Instructor Comments

Up until this week, we've dealt with technologies used within a LAN. But almost every LAN ends up connected to other networks, and that's the purpose of a WAN.

You connect to other networks by first connecting to a network belonging to an internet service provider (ISP). Through the ISP's network you can access the internet and you can connect to other LANs controlled by your organization – for instance, you can connect branch offices to a main office.

Connecting to the Internet

When connecting to the internet, the technologies used by businesses are largely the same that are available to residential users, though not all technologies are available in all locations. What makes business service different is that businesses often purchase guaranteed bandwidth and much improved customer service.

For residential use, you internet services are unlikely to guarantee any minimum bandwidth. Advertised bandwidths are approximations at best and often represent the maximum available. You are almost always sharing bandwidth with your neighbors. And if something goes wrong, you aren't going to be a particularly high priority.

But downtime for a business means lost productivity, lost sales, and lost customers. So they will often pay significantly more than residential users, but receive guaranteed service and responsive customer support.

Note this is not true of all businesses. The business service that you see advertised with residential service at a slight premium usually doesn't provide either of these benefits, it may be a repackaging of residential service that provides additional email addresses and a domain name, for instance. The services I am referring to are usually sold completely separate from residential service.

For organizations that are serious about avoiding downtime, it is common to have two ISPs, so that if one fails, the other can take over.

Connecting two LANs

The other use for WAN connections is to connect two LANs, for instance, connecting a branch office to the main office.

Obviously, if both the branch office and main office are connected to the internet, they can communicate with each other across the internet. But the real goal of a connection between LANs is to make the two LANs function as if they are a single LAN.

For instance, IP addresses used in LANs are typically private IP addresses, such as 192.168.n.n addresses. For a device inside the branch office LAN to communicate with a device inside the main office LAN across the internet, Network Address Translation (NAT) is required on both ends to translate the private IP addresses into public IP address. This is inconvenient and prone to configuration error. And then there's the security issue of having internal communications pass across the internet.

The types of connections that this chapter talks about allows the devices in the branch and main office to communicate with each other using the private IP address, no need for NAT. Routing protocols can even function across multiple LANs as if they are a single LAN.

WAN Essentials

I have nothing to add to this section.

WAN Topologies

This section is talking only about connections between LANs for an organization that has multiple locations.

You might think that the obvious thing to do is to connect every LAN to every other LAN. This is called a full-mesh topology.

The problem with a full-mesh topology is that it is very expensive – every connection between two LANs is an additional expense.

Consider for instance a chain of retail stores. You will want a connection from each store to the main office, but probably do not need a connection from each store to every other store, since the individual stores rarely need to communicate directly. But there may be five warehouse locations that do need to communicate directly, so there will be a full-mesh among those five locations.

This section just provides some examples how you may choose to connect the various LANs.

PSTN (Public Switched Telephone Network)

The next several sections describe various technologies used both within WANs and to connect LANs to WANs.

T-Carriers

I have nothing to add to this section.

Frame Relay

The textbook was good about describing the PSTN technologies are “legacy” technologies. Frame relay is not quite as “legacy”, but it is also a dying technology, most directly replaced by MPLS (below).

One other thing you should know about frame relay is that unlike most of the other technologies in this chapter, Frame Relay is NOT a physical layer technology, it can run over various physical media.

DSL (Digital Subscriber Line)

I have nothing to add to this section.

Broadband Cable

I have nothing to add to this section.

ATM (Asynchronous Transfer Mode)

The textbook probably gives ATM more coverage than it is worth, it is another dying technology that never caught on very much in the first place.

SONET (Synchronous Optical Network)

I have nothing to add to this section.

MPLS (Multiprotocol Label Switching)

For some reason, MPLS was introduced in Chapter 9. You should look back at page 464 at this time.

Also, like Frame Relay, MPLS is not a physical layer technology.

Metro Ethernet

Unlike all of the other technologies described in this chapter, Metro Ethernet is more of a concept than a specific technology. The concept is that to a customer, the connection to the ISP acts like any other Ethernet connection within their LAN. On the ISPs side of the connection there may actually be an Ethernet network. Or it might be that the difference between Metro Ethernet and another type of customer connection is that the transition from Ethernet to another technology just happens within the ISPs network instead of at the edge of the customer's network.

Wireless WANs

I have nothing to add to this section.

MISSING TOPIC – VIRTUAL PRIVATE NETWORKS (VPNs)

One of the topics skipped in Chapter 7 was Virtual Private Networks (VPNs). This is a topic you should know about. Chapter 7 goes into far more detail about cryptography than you need for an intro course, but this would be a good time to look back at page 340 and the description of site-to-site and client-to-site VPNs on that page.

In short, a VPN is a cryptographic "tunnel" between two devices, usually used over the internet. The name comes from the fact that the cryptography used protects the privacy of the data, and like the WAN technologies for connecting branch offices, it makes the two devices appear to be in the same network (hence "virtual network", yet another use of the term "virtual").

The reason for introducing VPNs with this chapter is that for smaller organizations, by far the most common way to connect LANs to each other is not through private WANs, but using site-to-site VPNs over the internet.

When you are paying for internet access already, the ongoing cost of a VPN is essentially zero. The technology is very likely part of routers the organization already owns, and if not, it is a relatively inexpensive addition. And the setup time isn't significant.

The downside of using a VPN is that if your internet connection goes down, so does your connection to the branch office. But this may be an acceptable tradeoff.

Troubleshooting WAN Issues

I have nothing to add to this section.