Chapter 9. Internet Application Software

Table of Contents

Introduction

Context

This unit follows on from unit 8, in that it extends and complements the contents of that unit.

Introduction

This unit gives a broad overview of the Internet's application software. It looks at electronic mail, file transfer, Telnet and the World Wide Web in terms of their operation, protocols, ways of ensuring uniformity, functionality with reference to communication and sharing, and utility.

All of these Internet applications are distributed: the meaning and implications of this statement are drawn out.

At this stage, half-way through the module, this unit reviews the situation that has been reached. The module has established the means of creating, and the idea of, a data communications fabric. This can cover a quite small area or can be global. It will be created in a different way according to its extent but will in consequence always ensure that the fabric is provided in an appropriate, affordable and effective manner. The module has also established that various types of application network can be built on this underlying data communications network. In other words, the underlying network enables various socially networked activities.

This can also be expressed in terms of layering. It has been shown that communication software can be built by adding to existing layers a new layer that, while it depend on the existing layers, also expands their capabilities. In the same way, an existing data communications fabric can, because it provides the basic means for communicating and sharing, support and be enhanced by an added layer of software that facilitates the activities of a community with members who need to interact with each other.

A final way of elaborating the same point is to indicate that a computer network makes possible both communication-enhanced computing and computer-enhanced communication.

Objectives

At the end of this module, you should be able to:

- understand in some detail the technical operation of the Internet's major applications;
- appreciate what is meant by a 'distributed application';
- construct an overall view of the data fabric provided by networks at all scales;
- appreciate the reasons for, and the benefits of, layering;
- explain the benefits of computer networking in terms of computer-enhanced communication and communication-enhanced computing.

Content

In parallel with this unit, you should read relevent chapters from you textbooks.

Introduction

The Internet has evolved through several stages in each of which a different application has been dominant.

The predecessor of the Internet, was the ARPANET. When it was established, it was expected that its main purpose would be to allow the sharing of the large and expensive computers that were attached to it. Although it was used in this way, its users, essentially the research community in the U.S.A., soon determined that it was of primary use as an e-mail network. Subsequently, with the widening of the user community, the most popular usage of the Internet has become the World Wide Web. So, in even its short existence, the major application of the Internet has changed from Telnet to e-mail to the World Wide Web.

In this unit, we will examine the operation of e-mail and the World Wide Web in some detail, including their specific protocols. The operation of Telnet and file transfer are not dissimilar. We saw at the end of the previous unit that the Internet provides a basic data transfer capability. All of the Internet's applications are provided in a way that takes advantage of this, and are, as it were, built on top of it.

To Do

Do Review Question 1.

The applications mentioned above are intended to meet the needs of the users of a computer network. To do this, they provide what may be see as generic services for communication and sharing. E-mail provides a form of computer-mediated communication for the users of computer networks. File transfer and the World Wide Web both allow anything that can be stored in a file to be shared. Telnet allows the computers on the network to be shared.

The issue of the uniformity of operation of these applications has to be addressed. Computers of many types are attached to the Internet: they can operate in different ways, and they can run different software for the same application. Despite this, the computers must all operate together when required to do so. This situation is attained by devising a standard way of operating for each application, and then creating programs that operate in this standard way for all the different types of computer.

To Do

Do Review Questions 2 and 3.

Electronic mail

Given the ability to exchange data over a network, it is not a large step conceptually to creating a message exchange capability and, from this to develop, as e-mail, an electronic version of the ordinary mail. The Internet's basic protocol for e-mail is known as SMTP.

We have seen that the text of a message can be coded character by character, using the ASCII code, for example, to convert it to data. By storing this data in a file at one computer and transferring it to another computer, a message has been sent by making use of the data transfer capability. This shows that an e-mail program is essentially a combination of a word processor, to create and store the file, and a file transfer program, to send it from one computer to another. It also shows that it is important for the computers involved in the exchange to use the same code for converting characters to data (imagine the consequences if they used different codes).

This simple conceptual view of e-mail demonstrates that it is an example of computer-mediated communication. A computer is needed to run the word processor to create a message, and a network is needed to provide the connectivity needed to communicate it to its destination.

This view also explains why it is straightforward to include attachments with e-mail messages. The file containing the e-mail and those containing the attachments are concatenated to form a single larger file which itself can be transferred from sender to receiver.

A technical view of e-mail can be given as follows. The numbered stages are keyed to the numbers on the accompanying diagram.

- The message sender uses the e-mail software to compose the message. This will be submitted to TCP for transmission to the local mail submission server.
- E-mail addresses. A destination address must be attached to the message. Addresses take the form 'mailbox@domain_name', for example, mit@cs.uct.ac.za.
- The mail submission server converts the domain name in the destination address to an IP address (using the domain name servers) and then submits the packets of the message to the Internet.
- The network routes the messages to the destination mail server.
- The destination mail server assembles the message and places it in the recipient's mailbox. The recipient's e-mail software can then display it.

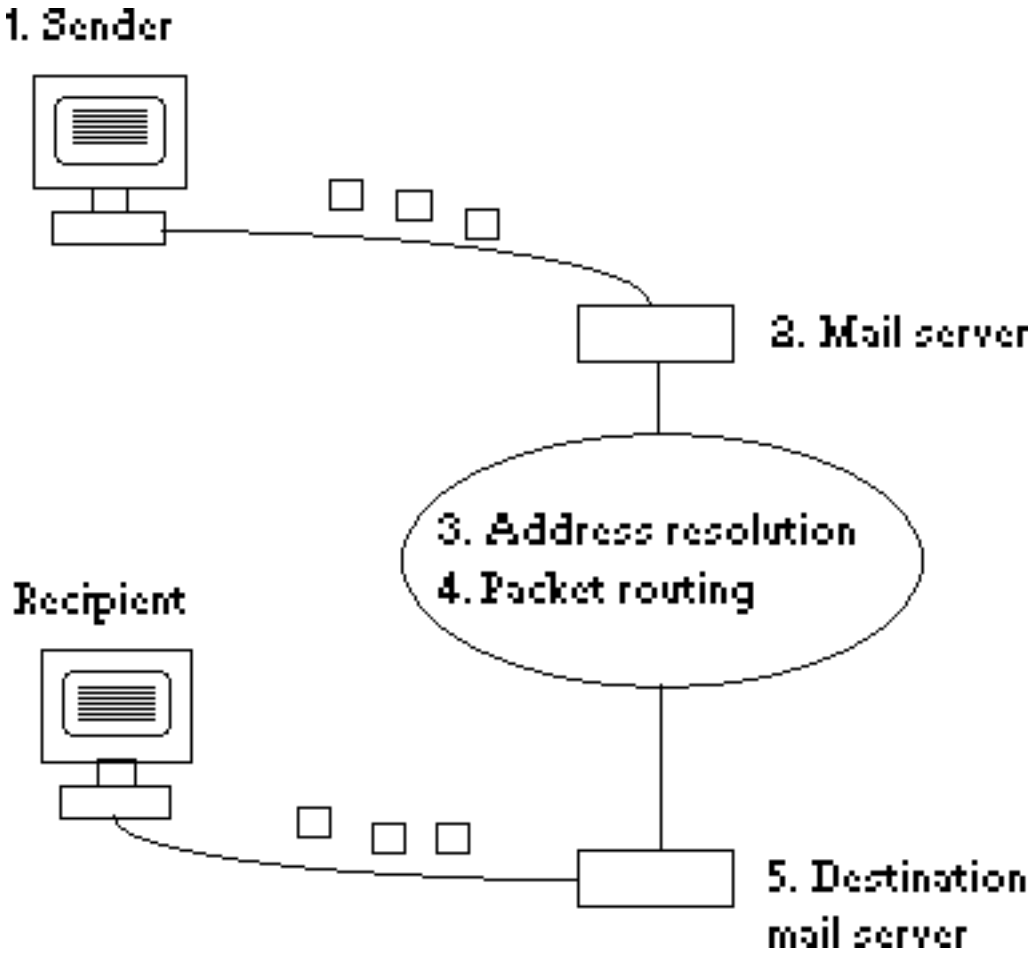

To Do

Do Review Question 4.

The World Wide Web

The information held by the World Wide Web is a distributed collection of information held by particular Internet computers that act as Web servers. This information can be accessed by any computer running a browser, part of the function of which is to turn it into a Web client. The pattern of operation between a client and a server is one instance of a well-known and widely used protocol known as the client-server protocol. This instance is known as HTTP.

The generic client-server interaction proceeds through four stages. They are listed below, and can be elaborated in the particular case of HTTP as follows:

1. Send request.

   In the case of HTTP, the request, as we have seen, takes a form such as:

   http://www.mdx.ac.uk/courses/index.html

   which is taken as a request to the computer named www.mdx.ac.uk for the file /courses/index.html'.

2. Receive request.

   The named computer (in this case, a Web server) receives the request and locates the file in its file store.

3. Send response.

   In this case, the Web server sends the requested file to the requesting computer. The file is of a particular type and assumes a well-known form.

4. Receive response.

   The client computer receives the file. In this case, having received it, one of the functions of the browser running on the client computer is to cause it to be displayed.

To emphasise that HTTP is a collaborative interaction between two separate computers, the programs running on the two computers are given below. They can be represented in the following schematic fashion:

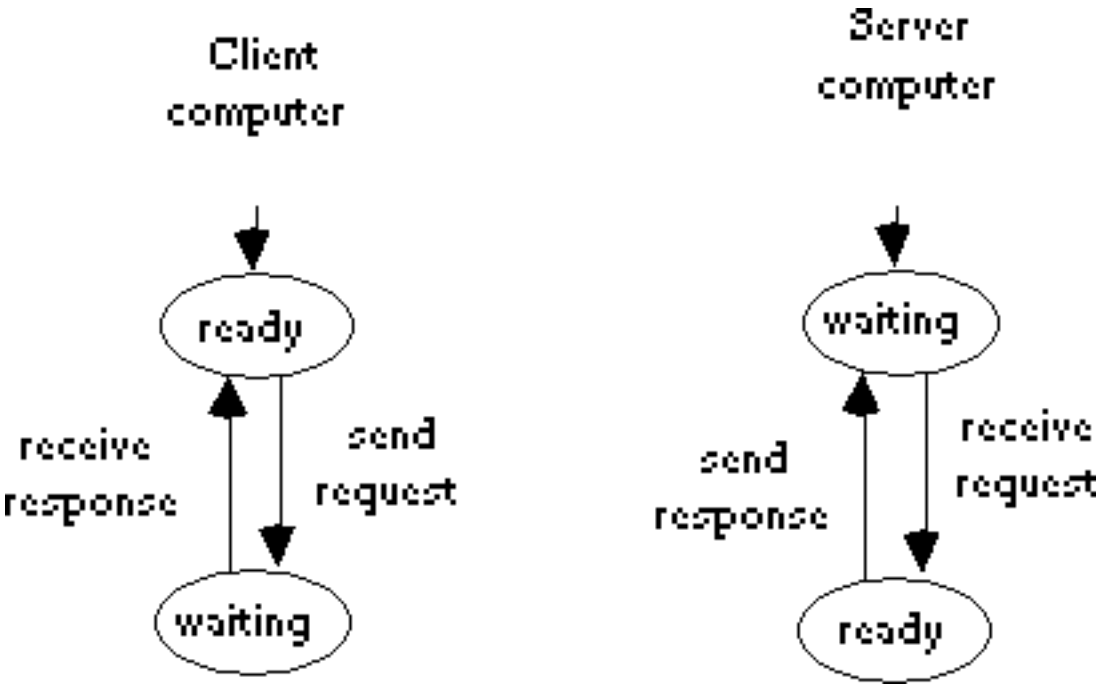

The computers can both be in one of two states, which are called 'ready' and 'waiting'. Initially, the client is ready to accept a request, and the server is waiting to receive a request. Obviously, the server cannot receive a request until one has been sent. Once a request has been entered at the client computer, it is sent, and the computer moves to its waiting state in which it is waiting for a response. The request will be delivered over the network to the server, which, on receiving it will move to its ready state in which it is ready to respond to the request. When it has located the file that was requested, it sends it to the client as its response, and moves to the waiting state in which it is waiting for another request (in general from another client). The response is delivered by the network to the client computer, which receives it and moves to its ready state in which it is ready to make another request (in general, to another server). At this point, the interaction is complete.

To Do

Do Review Question 5.

Distributed Applications

The applications discussed above are distributed in the sense that their complete performance depends on the co-operation of two separate computers that can be in quite different places. With e-mail, there is a sending computer at which the contents of the e-mail are composed and from which the e-mail is sent, and a receiving computer at which the e-mail is received and at which its contents can be read. On the World Wide Web, there is a client computer that requests a Web page and a server from which it

is to be obtained. In this way networked applications are quite different from stand-alone applications in that the computations needed to carry them out are distributed between different computers.

To Do

Do Review Question 6.

Review

It is worth reviewing the situation we have reached at this point in terms of understanding the ways in which a data communications fabric can be established, and how it may be established at any scale. We have seen that different types of networks are used at different scales of operation. On a small scale, local-area networks are used while, at the other extreme, the Internet has been created to provide a network with global coverage. In between are networks such as the national packet-switched networks and metropolitan-area networks. All of these networks, with the appropriate communications software, provide what we can call a data communications fabric, that is, a capability to move data transparently between any computers linked to the same network. The communications software is, as we have seen, vital to this and, in general its presence on all the computers of a network is taken for granted to the point that the computer and its communications software are treated as an inseparable unit.

The creation for a community, no matter what its scale may be, of a data communications fabric enables its members to exchange and share data with ease. This ramps up the community's infrastructure from a network which, although it provides connectivity needs to be understood at a technical level in order to be of any use, to a fabric that serves a basic purpose and needs only to be understood in terms of that purpose. Its ease of use will depend on the interface provided by the software. Its communication and sharing services can be further ramped up to support group working and co-operation within the community of users. The fact that data can be assigned any of a considerable number of meanings gives a data communication fabric considerable flexibility of application.

We have seen that the network supporting a community will be created in a different way according to the extent of the community: part of the reason for this is to ensure that the fabric is always provided not only in an effective manner but also in a way that is appropriate and affordable. To be effective, it must be able to support the various data-based applications and the socially networked activities needed by the community. Providing this in an appropriate and affordable manner is a matter for technical judgement.

To Do

Carry out Activity 1.

Layering

Layering is a key aspect in the construction of a data communications fabric. We can regard the physical aspects of a supporting network as a basic layer that is given increasing utility by the addition of layers of software to the point that it provides a data communications fabric. At this point, the addition of further layers can provide something more.

The first thing to say is that the communication process as it has been elaborated in the preceding units is complex, and the software that must be created to implement it is correspondingly complex. Some means of controlling, or organising, this complexity is necessary while creating the software. This has been done in various ways, which include the use of procedures and objects. With communications software, it has traditionally been done with layers. (In fact, a layer has a good deal in common with an object, and can be regarded as an object that can only communicate with the objects that correspond to its adjacent layers.)

The basic idea underlying the use of layering with communication software is that the addition of a layer to an existing layered structure creates a new structure with more capabilities than the previous

one. The added layer takes advantage of the existing capabilities to add something further to them. One theme running through the preceding units is that an existing structure can be used to support a new layer and that the new layer can enhance and expand the capabilities of the existing structure.

An important aspect of layering is that a layer, in effect, 'hides' the layers beneath it. This means that layering is a way of dealing with complexity not only in the construction of the software to deal with the communication process but also in managing the process itself. For example, a transport layer can be thought of as dealing with another transport layer. In fact, as shown in the diagram, it does not: it deals with its underlying layer which, in turn, deals with it. But this complexity is hidden. The underlying layers can be taken for granted, thus simplifying the way in which a complex situation can be conceived.

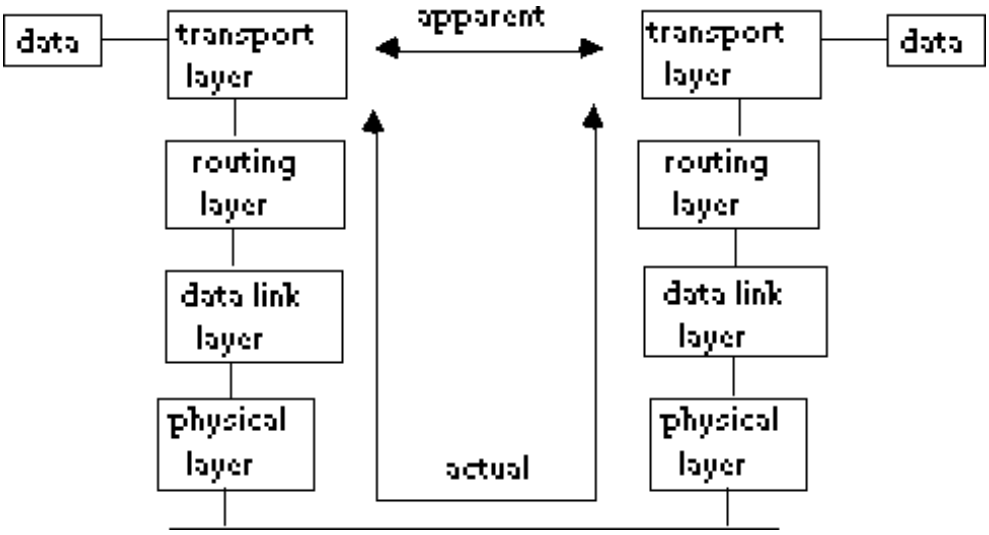

When the members of a community make use of their data communications fabric, they access it by means of one of the applications that are provided. These applications essentially belong to the application layer: this layer 'hides' the others, thereby providing transparent access to the communications fabric. The applications each provide their own user interface which is conceived in terms of the purpose of the application rather than the operation of the communications fabric.

Just as a data communications fabric can provide applications for communicating and sharing in its application layer so, it can support a further layer to link and co-ordinate these applications to provide higher-level services that enable co-operation in activities that require both communication and sharing.

To Do

Carry out Activity 2.

Synergy

The applications and services provided by a computer network can also be seen as a combination of the communication provided by the network and the computing provided by its computers. The combinations create synergy in that the combination is more than just the sum of the contributing parts.

We have, for example, characterised E-mail as a communication service, but we can also see it as an instance of computer-enhanced communication. Computers provide word-processing facilities for the creation and display of items of e-mail, as well as facilities for their storage, while the network provides for their delivery. In this way, the computing capabilities of a computer are harnessed to enhance what is essentially a communication activity.

Telnet provides for what can be seen as communication-enhanced computing in that it allows a networked computer to communicate with a larger computer, or a faster computer, or several other

computers so that its computing tasks may be carried out in some way that is better by harnessing the computing capabilities of other computers to which it is connected by a network. In this case, the computing capabilities of one computer are enhanced because it can communicate with other computers so as to gain access to and share their computational capabilities.

To Do

Do Review Question 7.

Activities>

Activity 1 - Requirements and Provision

In the previous chapters, networks of various scales have been described from what is essentially a technical point of view. The way in which they are constructed has been described as has the way in which, in consequence, they operate. It has then been shown that their mode of operation is suitable for the purpose for which they are intended.

Can you make the reverse argument? That is, by starting from the needs of a community of a particular scale, can you show that the solution that has been arrived at to provide them with a data communications fabric is, at least, acceptable, and perhaps better than that?

You can find a discussion of this activity at the end of the chapter.

Activity 2 - Layer Review

With the aid of your textbooks and other resources, review the development of communications software through the preceding units to draw out the way in which, starting from a physical network installation, the repeated addition of a new layer of software has not only built on the existing situation but has also increased its capabilities.

Activity 3 - The O.S.I. Seven Layer Model

Look up in your textbooks, the standard seven layer model for the organisation of communications software as agreed by the International Standards Organisation.

Confirm that the layers described in the previous units correspond to five of these layers but there is no equivalent for layer 6, the Presentation layer, and layer 5, the Session layer. Find out what the responsibilities of these layers are. Then explore the extent to which the matters for which these layers are responsible are needed in most Internet applications, and think about how their areas of responsibility are covered when they are needed.

Review Questions

Review Question 1

Describe in very general terms what has to be added to a data transfer capability to turn it into:

1. a message exchange capability (such as e-mail),
2. a file exchange capability (such as FTP),
3. a way of logging in to another computer (as can be done with Telnet), and
4. a way of requesting a page from a server (which is what the World Wide Web makes possible).

You can find an answer/comment for this review question at the end of the chapter.

Review Question 2

What is a protocol? Match these protocols and applications:

Protocols: SMTP, FTP, and HTTP.

Applications: file transfer, World Wide Web, and e-mail.

You can find an answer/comment for this review question at the end of the chapter.

Review Question 3

What is 'cross-platform program development'? What is its role within the realm of ensuring uniformity in the operation of applications?

You can find an answer/comment for this review question at the end of the chapter.

Review Question 4

Expand a little on the roles of address resolution, mail servers and TCP/IP in the operation of e-mail.

You can find an answer/comment for this review question at the end of the chapter.

Review Question 5

Verify that the client server interaction, as diagrammatically represented, moves from its initial configuration through the four stages described before returning to its initial configuration.

One way to do this is to represent the system configuration by the state of the client and the state of the server, and to track the state changes.

You can find an answer/comment for this review question at the end of the chapter.

Review Question 6

By describing the nature of the interactions involved in the operation of file transfer and Telnet, show that these are also distributed applications.

You can find an answer/comment for this review question at the end of the chapter.

Review Question 7

Characterise file transfer and the World Wide Web in terms of the synergy between computing and communication, explaining how the one enhances the other.

You can find an answer/comment for this review question at the end of the chapter.

Discussion Topics

1. Discuss the relative merits of:

   - creating a type of network for each scale of operation that may be perceived as being necessary, and

   - having a small repertoire of network types and a repertoire of devices for linking them (such as, for example, bridges and routers) so that a fabric of any scale can be created by linking existing network types.

2. It has been argued that layering is a bad thing. Can you think of any reasons for this? Could a fully object-oriented approach to the creation of communications software be preferable?

3. The Internet's communications software does not include Presentation and Session layers. Sometimes when communicating over the Internet, it is essential to provide for the matters for which these layers are responsible. How is it done?

4. To what extent do people in the following positions need to understand the technical aspects of the underlying network technology? Network users, Network service designers, Network service providers, Network managers, Network installers and Network designers.

Answers and Comments

Activity 1

The primary need of a community on any scale is connectivity. For a community whose members all operate in close proximity, a single cable to which they are all connected is the simplest solution imaginable. For communities whose members are scattered world-wide, the interconnection of existing networks is at least a pragmatic way, and perhaps as good a way as any, to provide world-wide connectivity.

The provision of appropriate support for connected communities then depends on the preparation and development of suitable distributed software that the interconnected computers can run.

Review Question 1

1. Taking it that a message is a sequence of characters, what is needed is a code that provides a representation consisting of a data sequence for each character.

2. Again, since a file is a sequence of items of some sort, what is needed is a code for the items.

3. What is needed is a code for the commands involved in logging in.

4. A code for the requests.

Review Question 2

A protocol is a set of rules governing the communications that can take place in a given situation. (It can also be defined as the set of messages that needs to be exchanged to achieve some goal.)

SMTP is the Internet's basic e-mail protocol, FTP is its file transfer protocol, and HTTP the protocol used by the World Wide Web.

Review Question 3

'Cross-platform program development' is the name given to the simultaneous development of versions of software for the same application that will run on different platforms (types of computer). Netscape, for example, used 'cross-platform program development' to develop simultaneously versions of their browser for the PC, Macintosh, Unix computers and so on. Following from this, the role of 'cross-platform program development' is to ensure uniformity of operation by developing software that behaves in the same way regardless of the platform (computer) on which it is run.

Review Question 4

Address resolution is the name for the process of changing the symbolic name for a computer, for example, www.mdx.ac.uk, to its IP address. The automation of this process is aided by the existence of address resolution servers.

Mail servers are servers that are used to submit the e-mails sent from a site and to store the e-mails delivered to a site. This simplifies the software needed by the computers used to send and receive mail because they need only communicate with their local server.

TCP/IP provides a communication service capable of delivering anything, including e-mails, so that the e-mail service can be created in a way that takes advantage of this existing capability.

Review Question 5

configuration	**Event**
Client ready, server waiting	
	send request
Client waiting, server waiting	
	receive request
Client waiting, server ready	
	receive response
Client ready, server waiting	

Review Question 6

File transfer involves the transfer of a file from one computer to another, so that the application is distributed between the computer requesting the file and the computer storing it.

Since Telnet provides a way to log into one computer from another, the application is clearly distributed between the controlling computer and the controlled computer.

Review Question 7

File transfer can be seen as an instance of communication-enhanced computing in that the ability of one computer to communicate with others that have a store of files makes it seem like a computer with a larger file store.

The World Wide Web can be viewed in similar fashion in that one computer is enabled to click through the files stored in the Web's distributed file store and to retrieve files from any Web server, making it seem, in both cases, as if it has a larger file store.