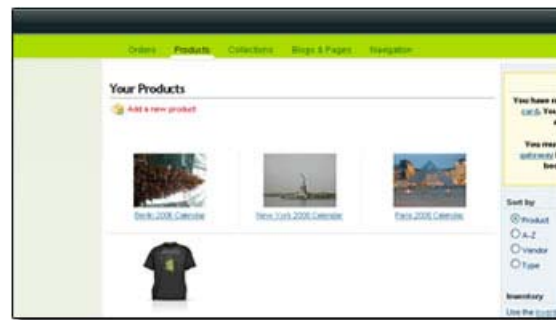


This material is taken from The Open University's OpenLearn website. OpenLearn provides free open educational resources for learners and educators around the world under a Creative Commons licence. For the online version of this unit and for other free educational resources across a range of topics, please go to www.open.ac.uk/openlearn

M360_1 An introduction to e-commerce and distributed applications



http://openlearn.open.ac.uk/course/view.php?name=M360_1

Introduction

This unit examines the type of system which is described by the umbrella term 'e-commerce'. A number of typical application areas are examined including retailing using the internet, supply chain management and online auctions. The unit also looks at some of the underlying technologies used to implement e-commerce applications, for example web technology. The final part of the unit looks at some of the problems which are encountered when developing distributed e-commerce systems, for example problems in ensuring that a system is kept secure from criminal activity. It concludes with an examination of a typical retailing system, how some of the technologies fit together and business models used in the internet.

Concepts

Anonymous remailer, B2B exchange, browser, checkout page, common gateway interface, cookie, day trading, denial of service attack, design pattern, disintermediation, distributed objects, dynamic pages, dynamic pricing, e-auction, e-learning, email server, e-mail, e-procurement, e-shop, e-tailing, file transfer protocol, framework, horizontal portal, hyperlink, hypertext mailer, Hypertext Markup Language, information brokerage, Java, online trading, portal, posting, procurement, query, rapid application development, search engine, Secure Sockets Layers, Server Side Includes, Servlet, shopping cart, spam, spider, stateless server, supply chain, third party marketplace, thread, trust brokerage, vertical portal, virtual community. web page, web server, website, webmaster.

Learning Outcomes

Having studied this unit you should be able to:

- detail what is meant by the term 'e-commerce';
- examine some typical distributed applications;
- detail some of the problems that are encountered when developing distributed applications;
- describe briefly some of the technologies that are used to support distributed applications;
- show how some of the technologies detailed in the unit are used in concert to realise a typical commercial system;
- describe some of the business models used in the internet.

1 Distributed systems

The past four years have seen an amazing growth of interest in distributed systems which address the business needs of companies and which use network technology – primarily the technology employed in the internet. In that period newspapers, magazines and even government pronouncements have frequently mentioned the huge business prospects for companies who embrace internet technologies. There is now a burgeoning literature on e-commerce aimed at the business person and the manager and a thriving industry of book publishing on specific technologies, for example Brown and Honeycutt, (1998); however, there is little published on how to combine the various technologies that are available to *design* and *implement* e-commerce systems.

Frankenedits

The internet floats on a mountain of jargon and acronyms. Frankenedit is an example of jargon. It is used by writers on internet topics and technologies who are often bullied or bribed by their publishers to write a book on specific technologies quickly in order to be the first to market. A frankenedit is what happens when such an author sends a book to a publisher for rapid editing and it returns in a worse condition than when it was sent.

2 Commerce and the internet

There are a number of ways in which companies can make money from the internet. Probably the best known way of making money is by selling some commodity; this could be a non-IT

commodity such as a CD or item of clothing or it could be some piece of application software, a font, a browser plug-in or an operating system. Other forms of revenue raising are:

- *Auction sites* which auction items on the internet and make profits by taking some commission from the sales.
- *Affiliate sites* which contain a link to a normal retailing site and are paid when a visitor from the affiliate site makes a visit to the retail site to make a purchase. The affiliate site will usually attract visitors by offering some information such as providing links to resources and tutorials on some specific topic or technology such as Java.
- *Banner adverts*. These adverts will contain links to the company doing the advertising; they will be displayed on a site and will result in some revenue being earned by the site owner when the banner advert is clicked.
- *Bulk-buying sites* where a site collects a number of users together all of whom want to buy some item; the site negotiates a discount with the supplier and takes a commission.
- *Shopping malls* where a number of e-commerce sellers congregate together on the same website; often these sellers will be related to each other, for example they may all sell luxury goods. The mall owner takes a percentage of their profit.
- *Portals* which contain massive amounts of material on a particular topic, for example a portal devoted to fishing. Such sites will contain thousands of resource links, tutorials and indexes. They will also contain links to merchants who sell goods associated with the portal topic. There may be a number of ways that the portal owner would make money, for example they could be paid by a merchant for each visit from the portal or the merchant may pay a flat fee for being included in the portal.
- *Digital publishing sites* which are effectively magazines on the web. They make profits in a number of ways including advertising and charging vendors for references to their website.
- *Licensing sites* which make some software available to other sites, for example search engines which allow a visitor to the site to search for material more easily.
- *Community sites*. These are like portals but involve the visitors more, for example a community site devoted to nurses might include a number of chat rooms which allow nurses to talk together in real time and swap advice. Money is made from such sites in the same way as with portals.
- *Name-your-price sites* are websites where the buyer haggles with the retailer and names what price they will pay for a particular product. Such sites make profits in the same way as normal retail sites.

Such applications have changed the face of retailing, for example the fast communication of the internet has made bulk buying sites feasible and popular and has given rise to a number of novel commercial models. The most popular model is one which involves a pyramid of services, ranging from those that are free, to those which are charged at a premium rate. For example, a site which sells a piece of software might give the basic software away for free and then offer increasingly more sophisticated versions of the software to buyers. This form of partially free charging has percolated down from the internet to conventional software sales; for example, the company Qualcomm that markets the *Eudora* email reader makes a version of the program available for no cost, but will charge for fully featured versions.

3 E-commerce applications

3.1 An example – supply chain management

Before looking at the wide variety of e-commerce application areas that have flourished over the last decade in more detail, it is worth looking at one which may not be familiar to a reader, but which saves companies huge amounts of resources. The application involves a **supply chain**. A supply chain is a set of relationships between a number of companies who have a symbiotic relationship with each other in that one company supplies commodities or services to other companies which, in turn, supply commodities or services to other companies, and so on.

The example was originally described by Kalakota and Robinson (1999) in their excellent management introduction to e-commerce. It concerns the processes involved in getting a bottle of Listerine mouthwash to the shelves of a retail chemist. It consist of the following steps:

- In Australia a farmer sells his or her eucalyptus crop to a processing company that extracts the eucalyptus oil from the leaves.
- The oil is then sold to a distributor in New Jersey.

- At the same time as the eucalyptus oil is being extracted natural gas is being drilled in the Saudi Arabian desert in order to produce the alcohol that is added to the raw ingredients of the mouthwash.
- Union Carbide ships the alcohol to Texas City, Texas, where the company that manufactures the mouthwash (Warner Lambert) has its factory.
- Farmers in the mid-west of America grow corn which is used in the manufacture of Sorbitol which both sweetens and adds bulk to the mouthwash. This is harvested and sent to the factory in Texas City.
- The ingredients are mixed and the mouthwash manufactured.
- The final bottles of mouthwash are sent to wholesalers or to the warehouses of chains of chemists from where they are distributed to individual retail outlets.

Figure 1 shows the information flows in this supply chain. This is an example of a **supply chain**. It represents a typical e-business application. There are a number of commercial imperatives for Warner Lambert: first, it should not overstock bottles of Listerine and incur costs because its resources are tied up in unsold goods; on the other hand, it should always stock enough bottles to satisfy demand in the time that the reordering and replenishing processes can take place. Second it should be able to get as quick a response to an order for raw materials as possible. This requires every company in the supply chain to move quickly to process an order from a company which follows it in the chain. An empty warehouse would be a disaster for Warner Lambert. A third commercial imperative is for the elimination of waste bureaucracy and indirect connections between companies from the supply chain. This is again connected with responsiveness: the more paper that is used and the more companies have to communicate using devices such as email the slower the process of reorder and replenishment will be. The ideal here is for a company higher up in the supply chain to share its data with companies further down the chain. For example, Warner Lambert should ideally be able to look at the stocks of the wholesalers and initiate a replenishment of those stocks when it discerns that they are becoming low – provided, of course, that some pre-agreement has been made about the quantities involved. An important point about an application such as this one is that information should be kept confidential as it flows across the internet.

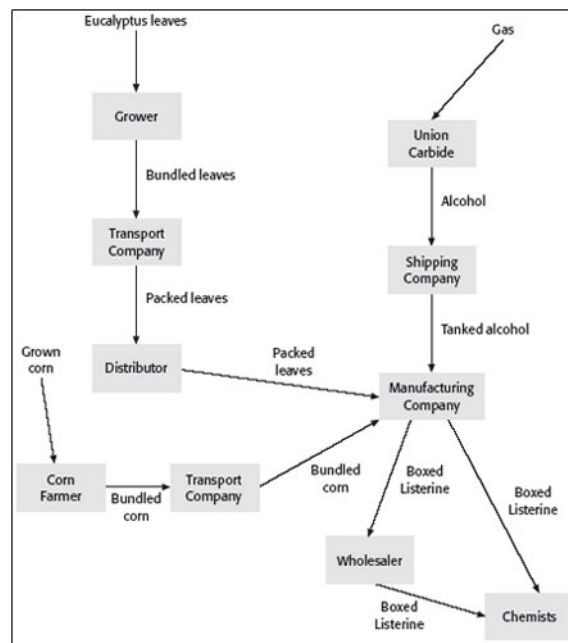


Figure 1: A typical supply chain

E-commerce and e-business

Internet terminology is still in a state of flux; nowhere is this more evident than in the past use of the terms *e-commerce* and *e-business*. Both have been used to describe any business activity which uses the internet. However, some consensus is emerging in that the terms are gradually being employed in a more focused way. The term *e-commerce* is increasingly being used to describe online retailing, for example the use of the web to sell books. The term *e-business* is increasingly being used to describe all business activities using the internet, not just online retailing.

All these imperatives lead to one conclusion: the fact that each of the entities involved in the process of manufacturing the mouthwash should be connected together. Even the farmer should be able to communicate with the buyer of eucalyptus leaves by typing in crop yields using, say, wireless technology.

Connecting companies together using network technology is not new: companies such as Procter & Gamble (The mouthwash example is an e-business application) have been leaders in this for years. However, the internet provides an infrastructure which enables individuals and companies to connect using a technology whose details are open to all and which can be easily implemented on virtually every computer in existence. The internet also provides a common interface to computing resources via a browser: everyone from the farmer in Australia to the manager of a retail chemist will be familiar with such technology and use it well.

A myth about e-commerce

One of the myths about e-commerce is that it is a comparatively recent phenomenon. Even in the early days of the internet when connections between individual computers were achieved by hand dialling using a telephone, there were a number of internet companies which had been set up to sell photographs and other graphic images of naked ladies and gentlemen. The earliest e-commerce applications were those associated with pornographers; indeed, a number of commentators have opined that the demands made upon the internet by pornographers have speeded the development of a number of technologies such as streaming video and the deployment of new business models.

3.2 E-tailing

The next example is probably the one that you expected me to introduce first: that of selling goods over the internet. However, I deliberately introduced supply chain management first since it is an area where companies are making huge savings in their investment in internet technology. The investments in retailing using the internet (**e-tailing**) are only gradually being realised.

The archetypal e-tailing application is that of a bookseller such as Amazon. This company is renowned for the fact that it only sells books over the internet and doesn't even take telephone orders. It has one of the best organised websites and is continually referred to by journalists as an e-commerce success story.

Customers of Amazon interact with its website and carry out a number of functions including:

- browsing readers' reviews of books;
- reading feature articles about books and authors similar to those found in magazines and newspapers;
- searching for details of a book based on information such as the author's name or the title of the book;
- browsing the books which are the Amazon bestsellers;
- ordering books using credit cards or some other similar payment method;
- tracking the progress of an order.

Behind the scenes of the Amazon site are a number of conventional functions which are found in all retailing applications, these include:

- *stock management*: keeping track of what books are in stock and ordering titles when stocks become low;
- *payment management*: paying suppliers of books for those that have been delivered;
- *customer payment management*: keeping track of payments made by customers and of payments made by credit card companies and banks which correspond to the customer payments;
- *delivery*: the process of sending books to customers;
- *market analysis*: the process of analysing sales in order to determine what books to order and which to discount in the future. This analysis occurs at both the customer level and at a temporal level in that customer preferences are processed and the times and dates when they express these preferences are analysed; for example, in order to answer questions such as what books sell well at Christmas or at Easter?

Most of these functions would be associated with any bookseller, irrespective of whether they use the internet or not.

Another myth about e-commerce

One of the myths about e-commerce is that the development of e-commerce systems is radically different from other commercial systems. I would say that it is *somewhat* different in that you have to worry about many of the problems that occur with distributing processing in a network; however, many of the functions required in the majority of e-commerce systems can be found in their conventional counterparts. Indeed, many e-commerce systems which are fronted by web servers still contain computers which were common ten years ago and are programmed in languages such as COBOL and C – languages which are not automatically associated with internet software development. Much of the analyses required for an e-commerce system are the same that you would carry out for a conventional system and also quite a lot of the design; however, they do differ in that the design of such systems is a lot trickier, for example to guarantee response times from a collection of computers communicating over the internet is a tough task.

3.3 Procurement

The term **procurement** is used to describe the purchase of goods and services which are not *directly* used in the main business of a company. For example, a car manufacturer will procure stationery for its employees or procure training courses for them to attend in order to improve their skills.

A typical conventional procurement process consists of a number of steps:

- the person making the procurement expresses their need by typing in details of a requisition using either a computer-based or paper-based form;
- the form is then dispatched to a member of staff who checks that it has been filled in correctly, that the amount is no larger than the amount that they are able to authorise and that there are sufficient funds available for purchase;
- if the form is authorised then it is sent on to a member of staff who is concerned with the purchasing of the good or service that is required; they then fill in a purchase requisition and send it off to the company who supplies the item that is to be purchased. If the item is over the limit for authorisation, then it is sent to someone who can authorise greater amounts.

This is in contrast to an e-procurement system which would automatically take the form produced by the person making the procurement, check that it satisfies all the company rules for procuring the item that is required, carry out authorisation if it is below a certain limit or send the form to someone who can carry out authorisation and then log the purchaser into the site of the supplier. He or she is then able to use this site to make the purchase, quoting an automatically generated procurement requisition number.

Again this is not hugely different to a conventional automated procurement system; however, it does cut out a number of inefficiencies at the purchase requisition end by virtue of the fact that the purchaser of a good or service is able to interact directly via the internet with the supplier. With procurement consuming as much as 10 per cent of a company's resources some large savings can be made by such an utilisation of e-commerce-based technology. Another example of the myth detailed in the previous section.

Microsoft and MSMarket

Microsoft discovered that 70 per cent of its purchases were for relatively small items which took up something of the order of 3 per cent of its purchase volume. The company discovered that a large amount of employee time was spent on the procurement process and hence invested \$1.1 m on a system known as *MSMarket*. When a Microsoft employee wishes to buy some item such as stationery they log into *MSMarket*, the system identifies them from their login identity and consults its database to discern what rules should be applied to purchases from that employee. The employee informs the system that they require some stationery and a screen of items and prices negotiated with a supplier are displayed. The employee purchases what is required and the

order is sent over the internet; an email is then sent to their manager to inform them of this and a tracking number generated which can be used to query the supplier if the item has not been delivered by a certain time. The use of *MSMarket* has increased exponentially since it was deployed and it now handles more than \$3 billion of orders.

3.4 Auction sites

These are sites on the web which run conventional auctions. There are two types of auction: those that are carried out in real time, where participants log in to an auction site using a browser at a specified time and bid for an article until the highest price is reached and no other bids are forthcoming. The other type of site – and the most common – is where an item is offered for sale and a date advertised after which no more bids are accepted. Such sites make a profit from two sources: first they usually charge a commission on the items that are sold and, second, they display adverts which are viewed by visitors to the site. The auction site will then receive some fee for displaying the advert, a further fee if a visitor clicks on an advert and it takes them to the advertiser's website and another fee if they purchase something from this site. Again, this is just an online analogue of a conventional business.

3.5 Other commercial websites

So far I have detailed e-commerce applications which are connected with very large organisations; to conclude this section it is worth looking at a number of smaller applications, many of which are distinguished by the fact that they are novel. They are in contrast to the applications discussed in previous subsections which mainly consist of standard functions such as order processing.

eCoverage and disintermediation

An increasingly popular e-commerce enterprise which does not involve e-tailing is that of quote finding. For example, a number of companies have set up websites which enable you to get quotes on various types of insurance policies including those for car, property and life insurance. One of the pioneers in this area is *eCoverage* who provide online car insurance quotes and who boast that it can insure you in minutes. This is an example of a growing trend in e-commerce applications: the cutting out of intermediaries; in the case of insurance this is the insurance broker. This is a process known as **disintermediation**.

3.5.1 Anonymous remailers

An **anonymous remailer** is a website which enables you to send an email anonymously to some recipient. The main reason for this is to do with something known as **spam**. This term describes unsolicited email which tries to sell the recipient something.

Spam

Throughout the internet you can find email addresses. They can be found embedded in web pages in the member's directories of internet service providers and in newsgroups. There are a number of companies who use programs known as **spiders** or **address harvesters** to search the internet for such addresses. These are then written to a file and sold to individuals and companies who then send bulk emails to the unlucky recipients. Often these emails are part of some crime such as selling bogus insurance policies. Spam is universally detested by internet users. Its name is derived from the Monty Python sketch which takes place in a café where a number of Vikings rampage round the café repeatedly shouting out the words spam, spam, spam. If you have not met the terms **internet service provider** and **newsgroup** before do not worry: they are introduced later in the unit.

You log into the site and type any message that you want to send; it will then forward the message on to its recipient with a dummy email address for the sender.

3.5.2 Link checking sites

The World Wide Web contains millions of web pages. Many of these pages are impossible to read, even though many existing web pages will reference them: your browser will usually return with some message such as 'Error 404 Page not Found' when you try to access them. Error 404 is a standard message returned by web servers when a non-existent page is accessed. It is also the telephone area code for Atlanta in the United States; you will occasionally hear technical staff referring to non-existent web documents as having 'gone to Atlanta'. There are two main reasons why a web document disappears from the World Wide Web: the first is that the developer or company might have deleted it, for example the company associated with the site has filed for bankruptcy or the individual who developed the site has moved it to another computer. The second reason is that the computer holding the web document is currently malfunctioning or has been switched off.

A link checking site is one to which you submit the address of a web page; it will store this address in a database and will then periodically check that the document is still accessible. If it discovers that a document is no longer available then it will email the customer who asked for the site to keep an eye on the document.

This is the type of service that technical staff who look after a collection of web documents find valuable; such staff, often known as **webmasters** (a term applied to both male and female staff, although the equivalent female term **webmistress** is very occasionally used), need to know very quickly when this happens. For example, the page that is no longer accessible could be the home page for a company that sells some goods through the internet: having web documents unavailable means that it effectively shuts the front door of the store to customers. Because speed in this case is essential a number of link checking sites offer webmasters a notification service via a customer's portable phone or pager.

3.5.3 Archive sites

These are websites which offer customers a facility for storing their files at a safe location. This guards against anything disastrous happening to the customer's computer and their losing valuable data. Often the files will be duplicated at a number of computers at different locations in order to guard against the possibility of one of the locations being affected by a natural disaster such as an earthquake, or a computer being affected by a catastrophic failure which results in its stored data being destroyed. The user of such a site usually registers with it using a name and a password; they are then presented with a set of instructions which take them through the process of collecting their files together to send to the remote location.

3.5.4 Change notification sites

These sites are a variation on link checking sites. Here, the customer is notified not when a web document becomes unavailable, but when the document is changed. For example, the customer might be interested in a particular page which advertises some holiday package offers to a particular destination and wants to keep abreast of any changes to the page which might signal the fact that a new improved offer has been added.

3.5.5 Email providers

These are sites which provide free email facilities; often they provide other facilities such as sending anonymous mail and constructing mailing lists. Such sites are valuable to users who are too impecunious to be able to afford conventional mailing software and to frequent travellers who can access such sites anywhere in the world. Their main disadvantage is that they tend to be slow compared with conventional mailing utilities such as *Microsoft Outlook* and *Eudora*.

B2B and B2C

The internet is awash with acronyms. Two acronyms used within e-commerce are B2B and B2C. The former stands for Business to (2) Business while the latter stands for Business to (2) Consumer. B2C is used to describe those business ventures which use internet technology to sell goods and services to internet users, for example the online selling of insurance policies is an

example of B2C. B2B describes the use of the internet for business transactions between companies, for example the holding of online auctions of bulk commodities such as crude oil. Current business thinking is that although B2C applications receive the most publicity it is the B2B applications which will have the biggest financial impact. There is also C2C (Company to Company) which is business between companies, for example the use of a network when two companies join together in some commercial activity such as building a shopping complex. The acronym C2C is occasionally used to describe ventures where consumers interact together, for example bulk buying sites; in this case it stands for Consumer to (2) Consumer.

3.6 Search engines

The web contains a huge amount of material. Finding specific information is a huge problem; even in the early days of the internet this was a problem which threatened to slow the growth of the net. Fortunately a partial solution to this problem emerged: the **search engine**. This is a program which accesses a huge database of information about the World Wide Web; it contains individual words in web documents and the location of the documents containing the words. When the user of a search engine wants to find any document they type a **query**: a series of keywords joined by Boolean connectives such as 'and' and 'or' or, in some cases, a natural language sentence. For example, the query

Java & compiler

would return with the addresses of all those web documents which contain the words 'Java' and 'compiler'.

In order to build up a search database the search engine will employ a program known as a **spider**. This will visit a website and access the web documents stored there, keep track of the address of the documents and the words that are stored in them and update the search engine's database. Spiders will not visit websites randomly: they will only visit those sites whose developers inform the search engine they want them linked to the engine's database. A developer will interact with the search engine site by requesting and filling in a form; this form will normally just ask for the address of the document to be indexed and a contact email address. After a few seconds the spider will visit the website and start the indexing process; usually, after a week or two, details of the website are added to the search engine's database. This is a description of those search engines which carry out automatic indexing. There are a few general search engines where the indexing is done manually by trained indexers. Two examples of this type of search engine are *Yahoo* and *Ask*.

Spamdexing

One of the tricks used by companies to make sure that their website is placed first when a search engine retrieves the results of a query is to include certain keywords a large number of times in the web documents in the site. For example, the word 'Java' repeated a large number of times will ensure that the site is displayed prominently when any search using this word takes place. This technique is known as spamdexing. Companies go to huge lengths to disguise spamdexing as search engine companies look upon the practice with huge disdain and will de-list any spammed pages. One technique that is used is to have a web page displayed with a graphic that has a coloured background and write the repeated words in the same layer and in the same font as the background.

Search engines are big business on the internet. They mainly make money by displaying banner adverts or sponsored links. There are a wide variety of search engines on the internet ranging from those which catalogue any website to specialised search engines which catalogue websites which address a single area such as Shakespearean studies or the LINUX operating system.

4 The facilities of the internet used to support e-commerce and e-business systems

4.1 The World Wide Web

The aim of Section 4 is to describe briefly the main facilities of the Internet that are used to support e-commerce and e-business systems.

The web is nothing more than a collection of files stored at locations throughout the world. These files are written using a special language known as the **Hypertext Markup Language** (HTML). A file written using this language will contain text which forms the information content of the file, together with instructions which define how the text is to be displayed; for example, HTML contains a facility whereby blocks of text are specified to be displayed as bullet points.

The user of the World Wide Web employs a program known as a **browser**. When the user wishes to read a file on the World Wide Web they will inform the browser of its address on the web and the browser will fetch the file. The browser will then examine the contents of the page and will determine from the HTML in the file how it is to be displayed; for example, it might meet some HTML which switches the display of the material from one font to another font.

A file which is downloaded into a browser is known as a **web page**. The computer that holds web pages is known as a **web server**. The collection of pages which are linked by some theme – for example, they may be pages which all belong to the same retail company – is known as a **website**.

Each page that is downloaded into a browser will have references to other pages expressed as **hyperlinks**. For example, a page belonging to a book retailer will have hyperlinks to the various sections of the site which deal with different types of books. Hyperlinks can refer to pages within the same site or can refer to pages within another site; for example, an online magazine might refer to other online magazines which are part of the same publisher's stable. Figure 2 shows a typical display from a browser. It represents a page from a site run by a British consumer organisation. There are hyperlinks embedded in the site in the main parts of the text (these are underlined) and hyperlinks in the left-hand side of the page in the shaded square.

The description above of the World Wide Web is a bare bones one which was true about eight years ago: web pages can now contain a wide variety of media including audio files, video files, graphics and even programs which can execute while the browser is being viewed. Without the World Wide Web e-commerce would be barely possible: it provides a standard interface to a variety of documents, products, services and software.



Figure 2: A typical web page

4.2 FTP

The acronym FTP stands for the **File Transfer Protocol**. It provides the facility whereby files can be downloaded into a computer from another computer in the internet. Although there are a number of utilities for file transfer most users now employ browsers for this via FTP links.

There are a number of utilities which enable you to load anything from clip art to the latest updates for operating systems. Many of these utilities are very primitive: they use a simple command line interface which lets you log in to the computer which holds the files, and then enables you to use simple textual commands to identify the files to be downloaded. However, there

are now a large number of sophisticated FTP programs which, for example, allow you graphically to show the structure of the file system on the remote computer, use drag and drop to download files and resume processing when transfer is interrupted by a network hang-up. web documents can also contain FTP links which also enable the downloading of files.

FTP is the mainstay of commercial companies who sell electronic products; it is a simple facility which has been found on the internet since its inception in the 1980s.

4.3 Email

This is one of the most ubiquitous technologies on the internet and, along with the World Wide Web, is the most used. When you write an email you use a program known as a **mailer**. When the email is completed it is sent via a number of computers known as **email servers** and via a number of other intermediate computers before it reaches its destination where it is read. In e-commerce applications email is a subsidiary, but important technology. It is used as the transport medium for mailing lists, for enabling customers to communicate with a company, for sending documents and data to customers and for keeping customers up to date about current products and services. Mailers are sometimes known as mail user agents while mail servers are sometimes known as mail transfer agents.

4.4 Newsgroups

A **newsgroup** is a collection of internet users who are interested in a particular topic. The topic may be a technical one, for example the LINUX operating system, or a recreational one such as fly fishing. Members of a newsgroup send messages associated with a particular issue such as the date of release of the next version of LINUX or the efficacy of using certain flies on certain rivers. Each message – known as a **posting** – will contain the user's thoughts on the topic. Once posted these thoughts are responded to by other users. For example, one user may say that they have got solid information that the next version of LINUX will be released next week. The collection of responses to a posting and the original posting is known as a **thread**.

Newsgroups can be moderated or unmoderated. If a newsgroup is moderated a member will examine each posting and determine whether it should be posted. There are a number of reasons why postings are rejected: one major reason is that it is not relevant to the area that a newsgroup covers; another reason is that the posting is abusive to another user. There are no restrictions on posting to unmoderated newsgroups.

Newsgroups are accessed by using a special purpose software utility known as a newsreader; although there is an excellent search site known as *Deja.com* which allows access to newsgroups. Figure 3 shows a particular posting displayed by *Deja.com*.

The September that never ended

In the early days of the internet newsgroups suffered a major drop in the quality of contributions every September. This was due to the fact that many students who commenced university at that time, were given their first internet accounts and started using newsgroups. After a comparatively short time, the standard of discussion rose as the students realised what was a valid and what was an invalid posting to a newsgroup. However, one September, according to internet veterans, the standard of contributions fell even further and never really regained its previous standard. This was the September when the large American internet provider America Online allowed its members access to newsgroups. This September was known as the September that never ended.



Figure 3: A typical posting displayed by Deja.com

Many newsgroups form part of the collection known as Usenet. This consists of a large number (30 000+) of newsgroups which are organised hierarchically and rooted in a number of categories

rec

such as

which designates newsgroups that deal with recreational topics and

comp

which deals with computing topics. The newsgroup designated as

comp.fonts

, shown in Figure 3, is an example of a technical newsgroup which forms part of the

comp

hierarchy.

Newsgroups are normally employed by ordinary users of the internet and have not really been associated with e-commerce. However, a number of companies are beginning to wake up to their potential. For example, a number of software companies assign staff to read the postings in newsgroups which are devoted to one of their products in order to field any questions which might arise about them: it provides a good impression to future and present customers if a company will provide help about a product without, for example, users having to ring a high-tariff phone line.

Other companies are also beginning to embed newsgroup technology into their web pages in order to create customer feedback groups which enable them to decide on future upgrades and new products.

4.5 Mailing lists

Mailing lists are groups of users who have some interest in common, for example they may all be network professionals. Such a list is used by organisations or individuals to inform the members of topics of interest to them. For example, my local cinema has a mailing list of cinema goers who have bought season tickets. It emails everyone on the list with the titles of those films which are to be shown in the coming week and notifies them of any special ticket offers. While there are many uses for mailing lists *within* companies there are also plenty of uses in e-commerce. For example, a mailing list can be used to inform current customers of any new products or services that are being offered. Most mailing lists are automatically maintained by specialised software. Such software allows someone to subscribe to a mailing list or drop out of a mailing list by just sending a simple email message to the software; for example, often all that is needed to subscribe to a mailing list is a single line email containing the message

Subscribe

This will result in the user who sent the email being added to the list of users associated with the mailing list.

Kelly's Rules

This unit is about the technical processes that are involved in the development of e-commerce and e-business systems. However, it is worth saying in passing that e-commerce and e-business applications seem to be radically changing the face of business. Probably the best chronicler of these changes is Kevin Kelly, one of the founders of *Wired* magazine. His most influential work is *New Rules for the New Economy* published by Fourth Estate. In this book he shows how e-commerce and e-business have overturned many of the conventional laws and rules about business. For example, he shows how companies can make huge profits by giving away free products, such as operating systems and browsers, with the profit being made from hardware, support software and services.

5 Issues and problems affecting internet, e-commerce and e-business development

5.1 Legacy technology

The aim of Section 5 is to examine some of the issues and problems which affect the development of Internet, e-commerce and e-business applications.

The World Wide Web was developed as a way of dispensing documentation within the large research laboratory at CERN in Geneva. I am sure that the originator of the technology, Tim Berners-Lee, did not realise at that stage how it would expand and become a major component of our economic infrastructure. Because many of the developers of the technology were unaware of its potential there are a number of problems associated with its huge expansion. Three of these are discussed in the following subsections.

5.1.1 Space problems

Probably the best known of these is the fact that the internet is running out of space for identifying computers. Each computer in a network needs to be identified by a unique data pattern known as an IP address. The current technology used to transport data around the internet is such that in the comparatively near future we shall run out of space to hold these unique addresses. Happily this is a problem that has been identified and groups of researchers around the globe have

developed new technologies which will eventually overcome this problem, one of these technologies being a new version of the protocol used to transfer data over the internet.

It is tempting to describe this in more detail; however, remember that this material is about developing applications for the internet: I shall confidently assume that all the work that has been carried out to cope with the explosive growth of the internet will be implemented 'under the bonnet' and developers will not need to know the details of this.

The great renaming

Periodically parts of the internet have to be reorganised. In the late 1980s the administrators of the newsgroups within Usenet decided to reorganise the naming conventions used for the newsgroups. They adopted a hierarchic convention so that topics would include sub-topics which, in turn, contained sub-sub-topics and so on. This was known as the great renaming.

Having again made the point about what is not relevant to this unit it is worth looking at something that is: web servers, and the fact that what they were designed for is not what they are now being expected to do.

5.1.2 Stateless servers

Web servers are what are known as **stateless servers**. What this means is that in their pure form they keep no memory of what has previously happened to them between requests; for example, when a request is processed by a web server for a page they have no direct knowledge about whether the page request was made by the same browser that asked for a previous page to be returned.

While this was not serious when web servers were being mainly used for dispensing documentation (their original use) it is a serious problem in e-commerce. One example of this is the **shopping cart**, or as it is known in the United Kingdom, the **shopping trolley**. When you visit an e-tailer and purchase goods you interact with a simulation of a shopping cart which keeps details of the goods that you have purchased. At the end of your interaction a web page, often known as a **checkout page**, will display the contents of the shopping cart and present you with the monetary total of your purchases. web servers as originally envisaged are unable to do this as they have no knowledge of any previous visit: they would not be able to remember the previous purchase.

(This is not the full story: there is some indirect knowledge. web servers will keep details of the accesses to their stored pages in a simple sequential file known as a **log** file which is used for marketing purposes and for optimising the web server.)

In the comparatively early days of the web this was seen to be a problem and a form of programming known as **Common Gateway Interface** programming was developed which enabled a web server to have a memory. There are a number of other, more recent technologies which have been developed to cope with this problem. The first is **cookies**; these are chunks of data which are stored on the computer running the web browser and which can be accessed by the browser throughout their interaction with a particular website. Such cookies can, for example, store the data associated with a shopping cart. Another technology used to store state is **servlets**; this is a technology, which employs cookies, and which is associated with Java; it enables the programmer to develop reusable code that can be plugged into a server and which keeps data persistently in the web server.

5.1.3 No dynamic web pages

Another example of a problem with web servers which arises from their original functionality is the fact that web pages were designed to be static: they were files which were stored on a computer and delivered in their stored form to anyone using a browser to access them. Many e-commerce and e-business applications require something much more dynamic, for example there are a number of financial service sites on the web which provide customers with up-to-date stock and share prices. These prices are stored on web pages and need to change very frequently – often every few seconds. There have been a number of add-on technologies that have been developed in

order to cope with this problem.

Saloman Smith Barney and paying for education

Saloman Smith Barney is one of the largest and innovative financial companies in the United States. Its website is an excellent example of a feature-packed site. One part of the site which requires dynamic pages allows the visitor to plan the funding of their son or daughter's education; it prompts the visitor with a number of financial questions and constructs web pages which provide details of the investment needed to pay college fees.

One early solution is something known as a **Server Side Include** in which parts of a web page are marked as being dynamic and, prior to their being sent to the browser, they are updated with data that has changed. Servlets are also used to produce dynamic pages, for example they can be programmed to return specific web pages to a browser containing content loaded in from a database. Another technology which has become very prominent over the last two years is known generically as **dynamic pages**. This is a more flexible version of Server Side Includes which allows the Java programmer to insert data into a web page at specified points on a real-time basis. There is also a Microsoft implementation of dynamic pages known as **active server pages**. There are a number of other technologies such as *mod_perl* and *php*.

5.2 Security and privacy

The internet is not a particularly secure place. There are two aspects to this: the first is that information is widely published throughout the internet which can be used for criminal and near-criminal activities. The second aspect is that since the internet is an open system, details of its underlying technologies are freely available to anybody. This means that the way data passes through the internet is in the public domain; the consequence of this is that, theoretically, anyone with the right tools can eavesdrop on data passing from one computer on the internet to another.

Share ramping, book ramping and painting the tape

Internet application developers face security problems. This is not the only problem that faces internet users. The internet, and in particular the World Wide Web, has provided such a fast and anonymous means of communication that old forms of criminal activity have had a second breath of life. Share ramping is the process whereby rumours are started about a company which would result in its shares either rising or falling, for example a rumour about it being taken over. The criminals who started the rumour will then either buy the shares if they have fallen and make a profit when they rise or sell shares they had bought previously when the price rises. The internet makes communication so fast and anonymous that share ramping has become a major financial phenomenon during the last five years. Share ramping was once known as painting the tape; it is derived from the ticker tape machines which were used to communicate share prices to dealers before the 1960s. A less serious form of ramping has occurred on online book retailing sites which publish readers' reviews of books, where authors and the staff at publishers submit reviews under an assumed name and which greatly praise a book. This is known as book ramping.

It is worth examining the first problem. Already you have met one of the consequences of data being readily published on the internet: the fact that spammers can use programs known as address harvesters to send large quantities of unsolicited email to users. There are much more serious manifestations of this problem, for example a phenomenon that has occurred in the last three years is cyberstalking. This is where a user of the internet finds the details of another user's email account and harasses them electronically, sending them emails, contacting them via newsgroups and intruding into the chat rooms that they use.

The possession of an email address can even provide the means whereby someone can bring down part of a networked system. It is relatively easy to program a computer to send many thousands of emails to a computer which is handling email communication for a company or organisation; the volume of emails can be so high that the computer is unable to carry out its main function: that of enabling staff of the company or organisation to send and receive emails. This is a form of attack known as a **denial of service attack** or **degradation of service attack**. An example of this

occurred when Serbian nationalists flooded the main Nato email server during the attack on Serbia in 1999.

List linking

This is a recent form of harassment where someone discovers your email address(es) and subscribes you to a large number of mailing lists. Often these lists generate as many as a hundred emails a day and some also send emails with large file attachments associated with them. A malicious user who wishes to disable another user's email processing can easily do this by subscribing them to hundreds of mailing lists; this is a process that is quite easy to automate. An attacker who wants to disable the communications of a large company can, if they have access to the internal email directory of the company, disable its email system completely.

The second aspect of security is that data flow across the World Wide Web and the protocols used to communicate with computers in the internet are public. This means that anyone who wishes to enter a computer system which has a connection to the internet or anyone who wishes to read the data passing through it has a major advantage. There is, however, a contrary point of view which states that by keeping security details open any security breaches can be plugged easily by patches generated from a knowledgeable community of developers.

There are major gains for the criminal in being able to access a 'secure' system, for example a criminal who can read the details of a credit card passing along a transmission line from a browser to a web server, can use that data to order goods over the net and remain undetected until the next time the credit card statement is delivered to the card holder; in this respect they have a major advantage over the criminal who just steals the card. A criminal who wishes to sabotage a network – perhaps they are a disgruntled former employee of the company – can send a program over the internet which is then executed on the internal network of the company and deletes key files. A commercial spy can monitor the data being sent down a communication line and discover that it is from a company to a well-known research and development organisation which specialises in certain niche products. This information, even just the name of the R&D company, is valuable to any competitor.

How secure is the internet?

In 1996 Dan Farmer, one of the leading members of the internet security community, analysed a number of internet sites using a tool known as SATAN which reports on security vulnerabilities. He discovered that out of the 2200 sites he accessed, 1700 were relatively easy to attack (77 per cent of the sites). This is a staggering figure; however, what makes it more staggering is the fact that Farmer chose sites which should have been neurotic about security, for example sites owned by banks, government agencies, insurance companies and credit card companies.

When the internet and the World Wide Web were developed security was not high on the agenda. There were two reasons for this: the first is that the developers of the embryonic internet were tussling with what was then novel technology and most of their focus was on basic aims such as establishing and maintaining reliable communications; the second reason is that very few people realised then that the internet was going to be used for commercial purposes – a theme which the previous section detailed.

Happily there has been a huge increase in technologies used to secure the internet. For example, a technology known as **Secure Sockets Layer** uses cryptography to encode the data passing between a web browser and a web server so that anyone eavesdropping is unable to read it.

5.3 Programming and abstraction

In the early 1990s programming an application for the internet was a tough proposition. I remember that I once had an application which required a very simple form of communication with another application located at a remote computer. I used a technology known as Winsocks which required me to carry out some pretty arcane code development just to send a simple test message to another computer and to receive a reply from that computer.

Java, when it appeared in 1996, enabled developers to treat another computer on a network essentially as if it was an input or output device; the programming code required to send data or receive data from another computer differed only slightly from that required to send and receive data from files.

However, even the programming facilities provided in the initial releases of the Java system are in opposition to a principle that both the developer and the user of a networked system should be unaware of the fact that they are accessing a networked system. This principle has been enshrined in a sales statement from Sun Systems, the original developer of the Java language, that the 'network is the computer'. What this means is that the developer should be designing and programming in such a way that much of the detail of the internet is hidden away under a number of levels of abstraction.

This is best exemplified by the idea of **distributed objects**. Many of you reading this book will have used Java: you will have developed a number of classes which describe objects whose methods were executed in order to achieve some processing aim. The objects that you employed were all created and stored in the main memory of the computer which you were using to carry out the programming. I would hope that you found that the object-oriented paradigm was a useful one which allowed you to think of a system in an easily approachable way: as a collection of objects interacting with each other in order to implement some functionality.

Distributed objects are objects which are stored on computers in a network, and to which messages can be sent as if they were objects residing on the computer which is sending the messages. In this way a programmer develops software for a distributed system in the same way that they would for a single computer: by defining classes and by executing code containing objects defined by the classes, with the code sending messages to the objects; the actual details of how the transport of messages occurs would be hidden from the programmer. (This is not the full truth, somewhere someone still has to distribute the objects to the computers in which they are stored.)

Two distributed object technologies: RMI, which is a pure Java technology, and CORBA, which enables distributed objects programmed in different languages to interact with each other.

The theme of greater levels of abstraction does not stop there, however. A form of distributed programming known as tuple space development is as far from the physical details of the internet as you could possibly get. Here the underlying model is that of a large shared data space to which computers on a network can read and write data. In Java this is implemented using a specific technology known as *JavaSpaces*.

5.4 The speed of development

E-commerce consultants speak of a web year. This is the time which it takes to bring to implementation a conventional system that would normally take a calendar year to develop. Current estimates are that one calendar year is equivalent to seven web years. Nowhere is there more of an imperative for companies to develop products and services quickly, together with the computing infrastructure required to support them, than in e-commerce. In software engineering terms this has given rise to a number of software development methods which are loosely described by the term **rapid application development**. In technology terms it has given rise to a number of ideas which go some way along the path which ends with providing facilities that enable companies to develop systems by just bolting components together, with many of the components being specified using design templates.

The rapid development of object-oriented programming languages such as C++ and Java has meant that the last five years have seen a growth of technologies that enable a developer to program software components that can be reused time and time again in applications other than those which they were originally developed for. It is worth pointing out that the ideas detailed here are not just confined to the Java programming language components; patterns and frameworks could, for example, be equally applied to C++.

Over the last five years there has also been a major increase in interest in reusable design. In 1995 Erich Gamma, Richard Helm, Ralph Johnson and John Vlissides published a book which has become a major bestseller (Gamma et al., 1995). This book describes how many software systems contain commonly occurring **design patterns** and that these patterns could be catalogued and reused time and time again in the design process.

5.5 Structure and data

A problem that is being increasingly experienced by internet companies is the fact that they have to interchange a large amount of data and that such data inherently lacks structure. For example, HTML has proved to be an enduring markup language for developing web pages; however, there are no facilities within the language, for example, to indicate whether an item of data, say a three-digit number, represents the price of a commodity or some hourly rate charged by a company employee.

There are also problems with browsers. There are two main browsers employed by users of the World Wide Web: Internet Explorer and Netscape Navigator. Each of these browsers can display the browser pages they process in different ways, especially if they contain advanced facilities of HTML.

There is also a further problem with browsers which is even more serious than the one detailed in the previous paragraph. Networking technologies are now being used in conjunction with other technologies such as those associated with mobile phone technology and television. This has led to the emergence of a number of different markup languages which are focused on particular devices, for example there is a markup language known as WML (Wireless Markup Language) which is used to display documents on internet mobile phones. The diversity of such languages means that the overhead in maintaining a number of versions of a document for different media can be very large.

Happily a technology has been developed known as XML which can be used to indicate structure in a document. There are also a number of tools available which allow the developer to maintain a single version of a document expressed in a language defined by XML and easily convert it into a form that can be displayed on a variety of media including television sets, internet phones and a variety of World Wide Web browsers.

5.6 Problems with transactions

A distributed transaction is a sequence of operations applied to a number of distributed databases which form a single functional step. For example, a transaction which moves an amount of money from a customer's account to an account owned by the same customer is an example of a transaction. It consists of two operations: the operation of debiting one account and the operation of crediting another account. There are a number of problems associated with distributed transactions. This section will briefly concentrate on one. This is the problem of deadlock: the fact that a transaction applied at one server might be waiting for data which is currently contained on another server, with the other server awaiting some resource that is held on the first server. For example, the first server might contain the account data that the second server needs to complete a transaction, while the second server might require other account data for it to proceed. *Enterprise JavaBeans* removes from the programmer the need to worry about many of the problems detailed above. *Enterprise JavaBeans* is a distributed component technology which allows developers to develop reusable components which can be used in transactions.

5.7 Design

Designing a distributed system can also be a problem, for example the fact that computers in a distributed system are joined by communication media which can stretch over thousands of miles provides an added dimension to the design process in that response time can be a problem. Another, equally serious problem is that of reliability, for example the fact that a hardware malfunction can bring down a poorly-designed distributed system.

As an example of one design problem that a distributed systems developer has to face consider that of replicated data. Replicating data is such a common technique that there are a large number of products available that allow you to implement it without very much programming. A replicated database is a database which exists in the same form at a number of points in a distributed system. There are two reasons for having replicated databases: the first is reliability. When a system contains a number of replicated databases and one of them becomes unavailable – perhaps because of a hardware fault – another database can take over its role. The second reason is to improve response time. A designer of a distributed system will try and place a database close to its users, usually connected via a fast local area network. Often the original database that is used is a long distance away and can only be accessed via slow internet connections; hence replicating the

database and placing it close to the users usually results in a large reduction in response time.

However, using replication comes at a cost: each replicated database needs to keep up-to-date data and will need to coordinate with other databases in order to do this; this gives rise to synchronisation traffic over the network which supports the databases, and can result in a very slow response time. Designing for data replication, where the amount of replication and the location of the replicated data is such that response time is lowered, and yet traffic is not increased to the point where all the gains are nullified, is an art.

6 A distributed system

6.1 The architecture of a typical e-commerce system

Before finishing this unit it is worth looking at the architecture of a typical e-commerce system in order to see some of the technologies. This is followed by details of a real application which I shall use to discuss some of the issues involved in distributed system development. The description used is closely modelled on the Amazon site.

6.2 The application

The first application I shall describe is that of an online bookseller. Such a book sales system would carry out a number of functions:

- It would allow the user to browse through a catalogue of books.
- It would allow the user to browse through a list of the most popular books, with the list being updated every hour.
- It would provide the facility whereby a user can buy books and add them to a notional shopping basket.
- It would inform the user when the books ordered have been sent.
- It would provide customers with the facility to contribute reviews of the books that they have bought.
- It would send regular emails to users who have subscribed to an email list informing them of any recently published books and any special offers.

6.3 The architecture

The architecture of the system is shown in Figure 4. It consists of a number of components. The most important of these is the web server. This communicates with browsers used by customers.

There are two other computers that are used in the system which are directly connected to the web server: a mail server which sends and receives mail from customers and a mailing list server which administers the mailing lists of customers and their interests. Both these servers communicate with the web server, for example the mailing list server is periodically sent the email addresses of users who wish to be added to a mailing list. The system will have a number of backup servers, for example the web server will have a backup which is brought in if hardware problems affect the main web server. This, again, is an example of the second myth I detailed earlier: that e-commerce systems are totally new. Often such systems will contain quite old technology which is hidden behind internet-based technology.

The main data store for the system is a series of databases which contain details such as customer payments and the number of books on stock and on order. These databases are replicated and are stored on a separate computer which is a very powerful server. The package which maintains stock details and customer financial details is a specialised one which has been developed for booksellers; unfortunately it is not well suited for internet use since it was developed some time ago. In order for it to connect with the front end of the system – the part that is implemented via the web server – a number of distributed objects are stored on the server that holds the databases. The web server communicates its requirements by means of sending messages to these objects which then interact with the database software. The distributed objects act as a form of interface which hides the details of the databases and provides an object-based interface to the web server.

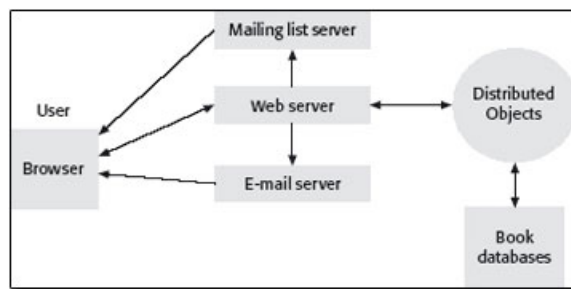


Figure 4: A simple e-commerce system

This, then, is a typical retail system that would use a number of the technologies and techniques: web servers, browsers, distributed objects, mail servers, replication and persistent state web server technologies.

6.4 The Sydney Olympic Games system

IBM was responsible for the computer systems which were used in the 2000 Olympic Games. There were a number of components to the system, these included:

- A website which was publicly accessible and which contained features on the Games, the competitors and the results.
- A Games management system which administered the logistics of the Games, for example arranging transportation, accreditation and accommodation for athletes.
- A Games results system which captured input from all the events in the Games and distributed them to judges, scoreboards, competitors, commentators and the website detailed in the first bullet point above.
- A commentator information system which provided real-time information to journalists and broadcasters, for example this system would flash up on a commentator's PC the times achieved by the runners in a race, only a few seconds after the race was completed.

The statistics associated with the development were staggering:

- 7300 PCs were used as clients.
- 2000 information workstations and kiosks were spread around the Olympic site.
- 1500 IBM staff worked on the system.
- 10 km of cable were laid.
- 815 network switches were employed to transport and send data.
- 600 servers were used to store data and provide other services.
- 13 million lines of code were written.
- During the 16 days of the Olympics the website had 230 million page views, 8.7 million unique visitors and 11.3 billion hits.
- The highest number of hits in a day for the website was 874.5 million.
- 371,654 email messages were sent to competitors from fans.
- The system achieved 100 per cent availability.

Clearly the project was a major challenge in software engineering, project management and logistics terms. It was also a major problem in terms of distributed systems development.

First, very high reliability was required. If a computer malfunctioned then it should not affect the functioning of the system, for example if the computer tracking athletes' timings in a race malfunctioned then the system should substitute another computer for it on-the-fly with no discernible difference to the users of the timing data.

Second, a large number of disparate pieces of hardware were used – both computers and output devices such as scoreboards. The system was developed as a classical client–server system in such a way that new hardware could be easily added. This was achieved via the use of standard protocols.

Third, high performance was required, for example results from the sailing events should have been sent to officials, journalists and competitors a few seconds after a race was completed. IBM carried out a large number of performance studies to ensure this was achieved and used techniques such as adaptive switching and data fragmentation to achieve this.

Fourth, scalability needed to be built into the system. IBM had made major investments in its

Olympics system and a main aim was that it should be capable of being reused time and time again, even if the number of sports, number of competitors and duration of the games get larger. Using a client–server system ensured that there is a high probability of this happening in the future.

7 Internet business models

7.1 What is a business model?

The aim of this section is to look at some of the business models which have been used to drive internet applications. A **business model** is a high-level description of an application type which contains all the common features which can be found in specific examples of the model. For example, one of the most popular business models is the e-shop which describes a website that sells products. The model is general in that it does not describe the item that is sold or the mechanisms that are used to carry out the sales process. The remainder of the unit describes a series of e-commerce and e-business models.

7.2 Business models

7.2.1 E-shop

This is the most ubiquitous form of commerce on the World Wide Web. It involves a company presenting a catalogue of its wares to internet users and providing facilities whereby such customers can purchase these products. Almost invariably such a site will contain facilities for ordering and paying for products by means of credit cards. The sophistication of sites described by this business model range from just the simple presentation of a static catalogue to the presentation of an interactive catalogue, the display of samples of products – for example the use of sound clips in a site selling CDs – the maintenance of mailing lists and the ability for customers to post reviews or customer reactions to specific products. This model is sometimes known as the storefront model.

Sites described by the e-shop model provide global presence, a cheap way to place products in front of an audience and decrease marketing and promotion costs.

7.2.2 E-auction

This model describes sites which electronically simulate the bidding process in a conventional physical auction. Such sites can range in sophistication from those which present a simple catalogue of items to those which offer multimedia presentations. Most sites which are described by this business model are concerned with selling items to individual consumers. However, there are an increasing number of sites which provide facilities for businesses to auction products to other businesses.

Revenues are raised by this form of site by charging for a transaction and for advertising. Some sites also sell the technology they use to other sites.

7.2.3 E-procurement

‘Procurement’ is the term used to describe the tendering of goods and services: a company decides that it requires some goods, say a fleet of cars for its salesforce. It would then announce this publicly and invite a number of auto companies to bid for the business.

Many companies are now switching to the web for the procurement process. A website devoted to procurement will normally advertise current procurement opportunities, provide forms facilities for companies interested in tendering and provide facilities whereby the progress of a tender can be tracked.

There are a number of advantages in carrying out the procurement process electronically. For suppliers it means that there are often more tendering opportunities, lowered cost of tender

submission and collaborative tendering with other companies. For the company offering tenders there is a major reduction in costs.

7.2.4 E-mall

An electronic mall or e-mall is a collection of e-shops which are often devoted to a specific service or product, for example an e-mall might be devoted to selling goods associated with a leisure activity such as fishing. Usually e-malls are organised by a company which charges the e-shops for administering their presence: maintaining the website, hosting the e-mall, and providing payment and transaction facilities and marketing.

The e-mall operator gains revenue for charging the e-shops; the individual e-shops have the benefits normally associated with e-shops, plus the fact that they are clustered together with other shops which operate in the same market segment and hence attract customers who might be browsing from shop to shop.

7.2.5 Virtual communities

A virtual community is a website which sells some product or service. In this respect there is no difference from an e-shop. The feature which distinguishes a virtual community is that the operator of the website provides facilities whereby the customers for a product or a service interact with each other, for example by pointing out ways a product can be improved. Technologies used for this interaction include mailing lists, bulletin boards and FAQ lists. The theory behind virtual communities is that they build customer loyalty and enable the company running the website to receive large amounts of feedback on the product or service they sell. A typical company that might run a virtual community would be a software supplier. Customers for software products manufactured by the company might post bug reports, bug fixes and work-arounds on a set of FAQ pages. Staff from the company would participate in the bulletin boards and also organise the FAQ lists.

Customers are often attracted to companies associated with virtual communities, particularly those that are maintained by companies that sell complex products, in that they see them as readily accessible stores of experience and unbiased advice.

A company can make profits from virtual communities in a number of ways. They can charge for participation in the community, and they can benefit from increased sales to customers attracted by the knowledge base held by the company and from a reduction in support costs.

The virtual community model is usually associated with another internet business model, for example the Amazon website is primarily an e-shop; however, the fact that it contains facilities for users to submit reviews and questions to authors and artists gives it the flavour of a virtual community.

7.2.6 Third party marketplaces

A third party marketplace is characterised by websites which offer access to a number of related companies, for example companies that are wholesalers of office stationery. A distinguishing feature of this model is that the companies delegate the marketing and sales of their products to the company that administers the marketplace. Typically a website which operates as a third party marketplace would provide a common interface to the products or services which are being sold, together with facilities for payment and delivery.

A third party marketplace is similar in some ways to the e-mall. The main difference is the fact that the product or service providers within the marketplace are more closely integrated, for example by virtue of the fact that there is a common catalogue interface to the products or services offered.

7.2.7 Information brokerage

Websites described by this business model offer access to information – usually business information. For example, a website which offers the results of surveys of customer satisfaction for a product such as a car would be used by car hire companies, auto companies and consumer organisations. Major providers in this area provide information derived from financial data such

as company performance figures, pension fund performance figures and financial market trends such as the growth of different types of mortgage. Companies whose internet presence can be described by this business model usually raise revenues by subscription or by a per-transaction charge.

7.2.8 Trust brokerage

This business model describes those companies or organisations who provide some service connected with security or trust. For example, as you will see later in the book, copyright is a major issue for the internet. A company might develop a sophisticated graphic which could easily be copied by another company that would then claim that they developed the graphic. A trust company might offer the facility for companies to register their work with them and then be able to testify to the date that the work was registered. Other trust brokers are associated with computer security and, for example, certify that a particular website run by a company is in fact associated with that company.

7.2.9 Collaboration platforms

Companies whose internet offerings can be described by this business model provide sites which enable companies to collaborate with each other, usually when the companies are spread over large distances. For example, a company which runs a collaboration platform might provide facilities for companies who wish to come together in order to tender for a complex project in a particular market sector such as aerospace.

7.2.10 Portals

A portal is a website which collects catalogues and characterises a huge amount of information. By displaying a large number of hyperlinks such sites provide an entrance (or portal) to the World Wide Web. Search engines originally provided fairly basic searching facilities; however, the past two years has seen them evolve *Portals on the nncmc* into portals. Portals are categorised as either horizontal portals or vertical portals. A **vertical portal** offers an entrance to large amount of information into a particular topic area, for example American football. A **horizontal portal** offers information over a large area.

7.2.11 Dynamic pricing

The dynamic pricing model is one which has a number of different instantiations. Basically, such models treat the price of a product or service (primarily a product) as variable and open to negotiation.

The **name-your-price** instantiation of this model is where the customer of a site offers the price that he or she thinks is reasonable for a product or service. The administrator of the website will pass on this bid to the provider of the product or service who will decide whether to accept it.

The **comparison pricing** sub-model encompasses websites which provide an interface to e-shops that sell a specific product. The model provides the facility for the customer to interrogate a database of product catalogues to look for the cheapest price for a particular product such as book or a CD.

The **demand sensitive pricing** sub-model is based on the fact that suppliers of a product will lower the price of a product if a number of units of that product are included in a single sale. websites which employ this model provide facilities whereby consumers can notify each other of their interest in buying a particular product such as a freezer. The site keeps a database of current products that have attracted a number of buyers with a predicted price and allow users to join the database of buyers who are committed to a sale.

The **bartering** sub-model allows consumers to barter services or products for other services or products. A site devoted to this form of economic activity will keep a structured database of items for sale and allows a buyer to barter with a seller.

7.2.12 B2B exchanges

A B2B exchange is a website or collection of websites which make the process of carrying out business to business transactions much easier. Under this banner comes sites which enable multiple companies to procure services and products from each other; help businesses form temporary alliances to carry out activities such as joint marketing or project bidding, and enable a marketplace in raw materials to function.

7.2.13 Online trading

This business model encompasses the trading of financial instruments such as bonds and stocks via the internet. Online trading has been a feature of the financial industry for some time. However, it was carried out using internal networks. The internet has enabled the individual user to trade stocks and shares from home and has given rise to the term **day trading**.

7.2.14 E-learning

This term is used to describe companies or organisations who offer educational courses via the web. The quality and features found in sites which can be described by this business model can vary. At its simplest such sites offer students the ability to download conventional texts. More complex instantiations of the model offer the students facilities to read individual lessons, try out online multiple choice questions and experience simulations relevant to the topic being taught.

7.2.15 Free products and services

It might seem paradoxical to include sites which provide free products or services under the category of business models. Typical sites which come under this category include gaming sites where users can play computer games using their browser, sites which run free raffles and sites which offer free software.

Such sites do not earn any revenues from the products or services they offer; revenue is earned indirectly, for example by means of banner adverts or by receiving revenue from sites which you have to visit before experiencing a service or buying a product.

One of the largest free product areas is that of free software. Organisations in this area include those who raise revenues and those who do not. An example of a company in the former category is Red Hat. This is a company that provides free versions of the LINUX operating system (LINUX is a free variant of the venerable UNIX operating system). You can download LINUX from the Red Hat website and install it on your computer without paying a penny to the company. Red Hat raise their revenues through support, packaging distributions onto CDs and providing services to companies who employ LINUX for application development. Companies such as Red Hat are the analogue of those companies who sell a razor for little or no cost but make their profit from selling the razor blades.

There are a number of sites in the internet which do not make any money from issuing software. These are sites associated with Open Source development.

The Open Source movement

The Open Source movement is a phenomenon of the 1990s. It is truly a child of the internet because it is only the Net that provides the communicational infrastructure to enable it to succeed. A typical Open Source project involves a number of programmers deciding to collaborate on the development of a useful piece of software. Initially the software might have been developed by one of them with some minimal set of functions. Over time the software is enhanced and undergoes a series of builds. The software is made available to the general programming community at no charge, including its source code. Other developers are free to modify the source code and release it. The Open Source movement has generated some very popular software systems which have gained a large number of adherents. The two most popular are LINUX and the Apache web server.

8 Further reading

There are a large number of books that have been written on e-commerce, many of which are of varying quality. I have found three useful. The first is by Kalakota and Whinstone (1997). It is an excellent introduction to both the technologies and applications involved in electronic commerce. Kalakota has also written a book on electronic commerce which avoids many of the clichés and which concentrates on unglamorous areas such as supply chain automation (Kalakota, 1999). If you ignore the jargon and the management speak in this book, you will find it a useful introduction to the way in which computers will affect major economies in the next decade. If you are interested in the economics of e-commerce then a good treatment is a book written by Choi, Stahl and Whinstone (1997). If you wish to delve further into the underpinning technologies of distributed systems then Coulouris et al., (2001) is the best introduction and is an excellent complement to this unit. A good introduction to e-business has been written by Timmers (2000).

The e-commerce gold rush

The e-commerce boom came to an end. Companies which were funded on the basis of a good idea and nothing else came quickly to grief, and the whole area of distributed applications became tainted with the fallout. However, writing in 2002 it is clear that distributed applications are still being deployed at a fast rate and that those e-commerce companies who still survive are making quite healthy profits. Surveys of e-commerce applications and e-business applications predict a steady growth in the use of technology mainly in the e-business arena and in the development of systems for individual companies.

Do this

Now you have completed this unit, you might like to:

- Post a message to the unit forum.
- Review or add to your Learning Journal.
- Rate this unit.

Try this

You might also like to:

- Find out more about the related Open University course [<http://www3.open.ac.uk/courses/bin/p12.dll?C01M360>]
- Book a FlashMeeting to talk live with other learners
- Create a Knowledge Map to summarise this topic.

Acknowledgements

The following material is Proprietary (not subject to Creative Commons) and used under licence (see terms and conditions) [<http://openlearn.open.ac.uk/mod/resource/view.php?id=15>].

Grateful acknowledgement is made to the following for permission to reproduce material:

Ince, D. *Developing Internet Applications*, chapters 1 and 4, published by Pearson Education Limited in collaboration with The Open University, © Pearson Education Limited, 2002, 2003. This publication forms part of an Open University course M360 *Developing Distributed and E-commerce Applications*.

The following designations, which are mentioned in the unit, are trademarks or registered trademarks of the organisations whose names follow in brackets:

Enterprise JavaBeans, Java, JavaSpaces (Sun Microsystems Inc); Amazon (Amazon.co.uk [<http://www.amazon.co.uk>]); Apache (Apache Foundation); Internet Explorer, Microsoft Outlook (Microsoft®); Eudora* (QualComm Inc); Linux (Linux Torvalds) Listerine (Warner Lambert); Netscape Navigator (AmericaOnline Time Warner); UNIX (AT&T).

Unit Image

JW_0000: www.flickr.com/photos/jw_00000/158078142/
[http://www.flickr.com/photos/jw_00000/158078142/]

Figures

Figure 2: reprinted with the permission of *Which?*, published by the Consumers' Association, 2 Marylebone Road, London, NW1 4DF, for further information please phone 0800 252 100;

Figure 3: permission approved by Google Marketing Department, USA.

All other materials included in this unit are derived from content originated at the Open University.

Every effort has been made to contact copyright owners. If any have been inadvertently overlooked, the publishers will be pleased to make the necessary arrangements at the first opportunity.

References

Brown, M. and J. Honeycutt (1998) *Using HTML 4*. Indianapolis, IN: QUE.

Choi, S., D.O. Stahl and A.B. Whinstone (1997) *The Economics of Electronic Commerce*. Indianapolis, IN: Macmillan.

Coulouris, G., J. Dollimore and T. Kindberg (2001) *Distributed Systems Concepts and Design*. Harlow: Addison-Wesley.

Gamma, E., R. Helm, R. Johnson and J. Vlissides, (1995) *Design Patterns*. Reading, MA: Addison-Wesley.

Kalakota, R. (1999) *E-Business*. Reading, MA: Addison-Wesley.

Kalakota, R. and M. Robinson (1999) *e-Business, Roadmap for Success*. Reading, MA: Addison-Wesley.

Kalakota, R. and A.B. Whinstone (1997) *Electronic Commerce*. Reading, MA: Addison-Wesley.

Timmers, P. (2000) *Electronic Commerce*. Chichester: John Wiley.

To access this unit online and to participate with others using this material please visit www.open.ac.uk/openlearn or put 'OpenLearn' and 'M360_1' into a search engine

OpenLearn and the Creative Commons Licence

This work (whole or in part) may only be copied and used strictly in accordance with the terms and conditions of OpenLearn and the Creative Commons Licence(s).

This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 2.0 Licence.



<http://creativecommons.org/licenses/by-nc-sa/2.0/uk/>

The Open University, Walton Hall, Milton Keynes, MK7 6AA
