# Internet and Intranet Technology

Workshop preview paper

by

**Richard Bland**

**Workshop features**

1. An overview of Web terminology and components

    - How the Web works
    - Web protocols and standards
    - Security
    - Intranets
    - Software support
    - Conclusion

2. How to build a Website

    - Building Web pages from common MS Office documents
    - Demonstration
    - Advanced techniques - Forms, CGI applications

3. Insurance Websites

    - A tour of some industry favourites

A short paper covering section 1 is attached.

The views expressed in this paper are personal and do not necessarily reflect the views of any organisation with which the author is associated. Whilst the author has used his best endeavours to ensure accuracy, any person or organisation using this paper to make decisions should check the accuracy themselves and seek their own professional advice.

## An overview of Web terminology and components

### 1. How the Web works

In its most basic form, the Web enables a user with a Web browser to connect to Web servers and read documents from them. The documents may contain text, graphics, audio and video clips: the latest versions of Web browsers allow designers to present their documents in almost any style. An important feature of these documents is that they contain *HyperText* links: words, pieces of text and pictures may be attached to links which specify the location of another document. The links are usually identified by being underlined and in another colour: clicking on them with the mouse causes the browser to switch to the linked document. The linked document need not necessarily be on the same Web server, or for that matter in the same country. A user can follow a logical trail from document to document by clicking on HyperText links without ever worrying about which Web server is supplying the information. Although the Web browser does not itself have a search capability, many Web servers have search facilities so that users can locate a Web server with a suitable starting point for their session.

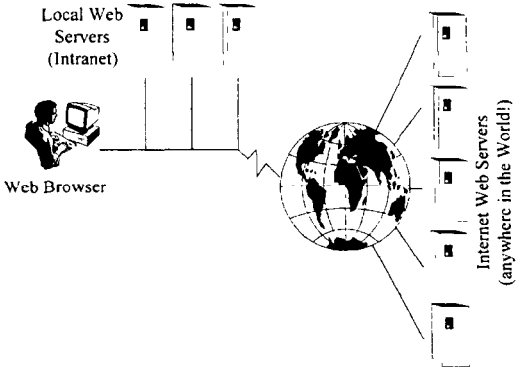


Web pages can be more than simple documents. Documents can be interactive: one possibility is the use of forms, which take input from the user and then supply it to an application on a Web or application server

for some form of processing. Another feature is the use of applets (tiny programs) which can be downloaded to the browser to perform some task locally.

## 2. Web protocols and standards

At the lowest level, all the nodes of the network are connected with TCP/IP network protocols: servers and clients are identified by unique four part addresses, e.g. 128.102.15.50, which uniquely define individual machines on the network. The network is tied together by routers which redirect data around the world network so that data packets flow between the two nodes which wish to communicate. There are many possible routes for data and no central hub: this was part of the design criteria for the original U.S. defence network so that it could withstand nuclear attack, but a consequence of this is that the network is impossible to police. It is not possible to prevent two nodes from communicating or to block users once they have an established a connection to the network, usually through an Internet service provider. Individual service providers can outlaw websites to their own users, but if a user can establish an independent connection to the web then they can go anywhere.

The next protocol is the HyperText Transfer Protocol: this defines the method by which Web browsers and servers exchange data. There are other protocols used on the web, e.g. FTP and SMTP, but the majority of websites and browsers rely on HTTP to communicate.

The documents used in websites are created using HyperText Markup Language. HTML is a relatively simple programming language which allows documents to contain text, graphics and, most importantly, hypertext links. HTML also allows the creation of forms which enable the user to supply information back to the Web server. The vast majority of Web documents are passive HTML which contain only hyperlinks to connect to other documents. However there are two ways in which the Web can become more active:

The use of the Common Gateway Interface in conjunction with HTML forms allows the user to activate an application program either on the Web server or on another machine connected to it. A typical application

might be an insurance quotation engine. The user fills in a form which is part of an HTML document on the browser and submits it to the server. The form identifies the application to be used, the server on which the application can be found, and supplies a range of data fields from the form to pass on to the application. The Web server uses the CGI to trigger the application, wherever it is found, and pass it the input data fields. The application then runs and generates output in the form of an HTML document which is passed back through the CGI to the Web server, which in turn passes it on to the browser. In this example, the output is probably a quotation complete with calculated premium.

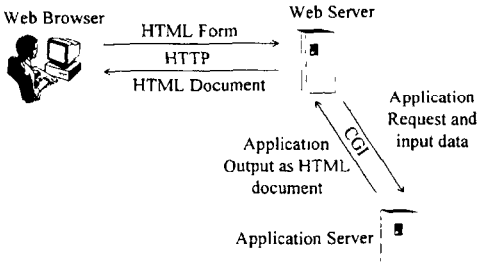


Alternatively, the browser itself can become more active by the use of Java applets. These are miniature programs which are downloaded from the server to the browser to endow it with extra features and capabilities. Java is the programming language developed for the creation of such programs.

## 3. Security

This is a problem for a network as open as the Internet. It must be assumed that a sufficiently determined hacker could intercept communications intended for any particular user. If a secure service requires a user to supply a userid and password, then it is possible that

without encryption these could be intercepted and used by an unauthorised person.

A solution to this is to encrypt the data passing across the network. Most efficient systems use symmetric key encryption, where both partners have the same key. But this creates the problem of getting the key into the hands of the remote partner. The solution is to start the session with asymmetric key encryption, using a public/private key pair.

Public/private keys are issued by a registration agency. The public key is supplied to all clients, and the private key is kept secure. The important feature of this key pair is that data encrypted by the public key can only be read using the private key, and data encrypted using the private key can only be decrypted with the public key. This means that:

(a) The client can send userid and password information to the provider using the public key, and only the provider can read it.
(b) The provider can send information to a client using the private key, and if the client can decrypt it using the public key then it must certainly have been sent by the provider.

Dual key encryption is very expensive in processing terms, but once the session has been established, a common session key can be agreed and then efficient symmetric encryption can be used for the rest of the session.

4. Intranets

An Intranet uses exactly the same technology and protocols as the Internet, but it is confined to a private network. Why use an Intranet?

It is much easier for users to find information on an Intranet than on a conventional network. For example, if a user wishes to find the phone book, the personnel holiday form or the latest quotation volume statistics, it will be much easier for them to follow the hypertext links in the internal company website than to remember a complex network reference.

In order to use an Intranet, the only client software required is a Web browser. The client PC needs no other product and can be configured with a minimal hardware specification. The use of CGI applications on the Web server may mean that quite complex transactions can be carried out, e.g. drill-down on a management statistics table, without using any client resource. The servers do all the work, and the browser just displays the results.

An Intranet is sometimes called an Extranet when it is extended to take in other companies, e.g. suppliers, clients, other group companies.

At this stage most Intranets are being configured to provide internal company statistics and standard internal management forms, e.g. stationery orders. But in theory they could be developed to provide a database transaction interface for internal clerks and external suppliers: in other words, an Intranet could be the normal clerical production environment of the future.

There is a big danger of reinventing the wheel, however. In the example described above it might prove much simpler to just have a normal terminal connection to some multi-user system running a database application rather than set up a complex series of Web and application servers.

## 5. Software support

Internet technologies are widely supported – if you have the latest versions of office software. Most of the recent releases of word processing, spreadsheet and database software packages have options to create output in HTML form. Microsoft Office97 products in particular conform to this specification and will be demonstrated as part of the workshop. The HTML standard is also appearing in many other products including the statistical packages favoured by actuaries.

The importance of this is that it becomes possible for actuaries and other statistical workers to produce their reports and output in a format which allows it to become directly attached to a website.

More interestingly, the CGI standard is also becoming available in easy to use ways. The original specification of this standard required the developer to write low-level programs in C or Fortran to direct the Web server to the application, to allow the application to pick up the input fields, and to format the output of the application as an HTML so that it could be supplied back to the browser.

Software suppliers, including those used widely by insurance companies, are now supplying custom built programs and add-ons which provide CGI capability to their applications without the need for complex programming. The module recently announced by SAS is typical of the approach: it consists of a broker program resident on the Web server which attaches to the CGI interface and directs calls to a SAS server. The SAS server has an application server process which picks up calls from the broker program and runs SAS programs in response to its requests. The point of such a configuration is that the developer has to develop the SAS application, but need know nothing of the complexities of CGI.

## 6. Conclusion

Web technologies are now mature. The software has been developed and is reliable: the communication technologies are tried and tested; and, perhaps most important of all, these are not systems which you have to go out and buy, but are supplied to you free when you buy your PC with all the latest office software.

From spreadsheets to generalised linear models, the IT world has produced a range of technologies which have become a way of life for actuaries. The Web is now one of them.