



**Poseidon House
Castle Park
Cambridge CB3 0RD
United Kingdom**

TELEPHONE:
INTERNATIONAL:
FAX:
E-MAIL:

**Cambridge (0223) 323010
+44 223 323010
+44 223 359779
apm@ansa.co.uk**

ANSA Phase III

Commercial Information Services in the World Wide Web

Nigel Edwards

Abstract

The ANSA Workshop on Commercial Information Services in Web Systems took place on Thursday May 12th 1994. The purpose of this workshop was to look at the problems and requirements of such systems.

This note summarizes the main issues which emerged from the workshop. It tries to capture the discussion rather than the contents of presentations or slides, since these have been (or we hope will be) made available separately.

The following issues were discussed: remuneration; resource control; recovery mechanisms and auditing; naming; caching; agents and brokers; usability; security; customer types; legal and accounting issues; CommerceNet; The World Wide Web Consortium.

APM.1220.00.12

Draft

20 May 1994

Briefing Note

Distribution:

Supersedes:

Superseded by:

1 Commercial Information Services in the World Wide Web

The ANSA Workshop on Commercial Information Services in Web Systems took place on Thursday May 12th 1994. The purpose of this workshop was to look at the problems and requirements of such systems.

This note summarizes the main issues which emerged from the workshop. It tries to capture the discussion rather than the contents of presentations or slides, since these have been (or we hope will be) made available separately.

1.1 Introduction and executive summary

Over recent months a dramatic take off of the World Wide Web (WWW) has occurred. Evidence of this is the huge number of new servers which are advertised each month and anecdotal evidence of the tremendous amount of internet traffic being generated by WWW. The main reason for this sudden take off is the accessible user interface and the cheap connectivity offered by the internet.

Even more recently there has been growing interest in exploiting this technology for commercial purposes (see §1.13). A growing number of commercial organisations are using WWW as a very cheap and effective way to promote their business. At present this tends to be (interactive) advertising¹, but there is growing pressure to conduct business transactions using this technology.

Thus the WWW is a driving force for new distributed applications which will increase the demand for client/server technology and networking. However, whilst the user interface is very good, the underlying platform and management technology is very old (it is essentially a development of the ftp and email technology). The pressure for commercial use is already showing up some of the problems associated with such technology. In addition there are difficult systems integration issues which arise when different organisations wish to collaborate to do business electronically. Many of these problems and issues were discussed at the workshop and are discussed in this paper; they are problems and issues which ANSA has been addressing for a number of years.

It was clear from the workshop that the WWW and internet technology is only the start, but it shows the way forward. In the future there will be a demand for services based on much richer kinds of (multi-)media to be delivered. There are fundamental reasons why these services cannot be delivered using internet technology. More modern technology such as ATM could deliver these

1. For example, DEC have two alpha machines which are accessible to users of WWW; at a recent meeting of CommerceNet the DEC representative reported that 26% of those who had used the Alpha in this way had bought it.

services, provided it can be controlled to deliver the right kind of performance guarantees.

Therefore commercial information services in Web-like systems are an excellent opportunity to use the ANSA concepts and technology. This application area can provide a framework in which to evaluate and develop the ANSA work on performance, dependability and federation.

The ANSA Workshop on Commercial Information Services in Web Systems confirmed the above hypotheses. The remainder of this paper is a technical summary of that workshop.

1.2 Issues Discussed

The following issues were discussed:

- Remuneration
- Resource Control
- Recovery mechanisms and auditing
- Naming
- Caching
- Agents and brokers
- Usability
- Security
- Customer types
- Legal and accounting issues
- CommerceNet
- The World Wide Web Consortium

1.3 Remuneration

There was much debate about the issues raised by the model for remuneration presented by Gray Girling, including just what exactly does it mean to charge in an electronic environment. Would it involve some kind of electronic currency or fund transfer; or would bartering be acceptable? Subscription may well be the most convenient form of payment for many services.

There is a need for lightweight billing mechanisms: if the information is worth a few pennies then a billing mechanisms which costs 50 pence per transaction is not suitable.

The need for one stop billing mechanisms was discussed. For example, a company receives a single bill for electricity (per site/entity) rather than one per employee. The same is desirable for certain kinds of information services. There may be a role for agencies to collect and distribute money, effectively providing a one stop billing service. However, there is also a need to control what users are doing; analogies were drawn with call-barring of certain telephone numbers.

Given the wide variety of possible remuneration strategies, fairly detailed information will be needed to support federation between systems with different strategies.

There was much discussion about the need for certification of electronic money. Each time electronic money was used in payment the parties involved would need to certify it to make sure it was genuine. The mechanisms put in place for this paradigm to work would need to be scalable: a central certification server would not work. The latter would also probably be unacceptable because there is unlikely to be a single global entity which everybody would be prepared to trust.

Payment must be auditable and non-repudiable if something goes wrong. A major defence against fraud is to distribute information and evidence into different places and organisations, so that you can audit it to check for inconsistencies.

Payment granularity affects the cost of providing a service. The relationship between charging mechanisms and cost is complex. On the one hand the charging mechanisms themselves require resources which tends to increase the cost of service delivery. On the other hand they can increase the effectiveness of service management which tends to reduce the cost of service delivery.

1.4 Resource Control

Resource control of the underlying network is a big problem for Web systems. The notion of Web trails was discussed: some trails would be busy motorways and require a lot of resources, others would be quiet backroads. Whether something was a motorway or a backroad would be time dependent. This may be helpful for predictive caching strategies.

Although the internet gets us some of the way, there are fundamental limitations which means it will not be able to deliver extremely time-sensitive services (e.g. interactive multi-media). ATM appears to offer more promise in this respect.

1.5 Recovery mechanisms and auditing

Underlying failure recovery mechanisms were a recurring theme. Business activities need to be dependable. They need to be auditable and traceable, so that inconsistencies can be detected and resolved. A notion of audit services was proposed as one way of doing this. There is a tension between this and the desire of some for anonymity.

Contracts and arbitration also have an important role to play in recovering from failure and in preventing failure. The former says what should happen or what should have happened; the latter attempts to resolve disagreements and disputes so that interaction can take place.

Notions of trade-marks and trade associations may be helpful when a business transaction goes wrong. The owner of the trade-mark or the trade-association would be responsible for correcting the failure or making reparations. There are security issues associated with being able to use the trade association label: is it genuine and can it be faked?

1.6 Naming and resource location

Naming was discussed in the context of what does it mean to name an object on the Web. There is a need to associate names with sets of objects: a set of objects contains zero or more related objects. For example all issues of CACM in 1993.

There is also a need to name objects which satisfy more dynamic attributes: for example the latest issue of CACM. The same issue occurs with the need to name a role rather than an individual, such as “webmaster”.

Naming schemes need to be scalable and federate (I would like to be able to choose my own naming scheme for objects I own). It needs to be scalable to 10^{15} objects or more.

1.7 Caching

This is very closely related to resource control. Scalable caching was identified as being an extremely useful piece of enabling technology for Web like systems to evolve and avoid bottlenecks. Caching can also be used as a cost control measure, by reducing an individual’s demand for network resources. There are consistency issues which arise when there are multiple copies of the same thing.

1.8 Agents and brokers

It is very easy to get drowned in information — it is hard to find the information you need. Agents are often discussed as a solution to this problem. There are at least two problems with agents:

1. How do users set up their agents: are they competent?
2. Even if you have an agent exploring the network for you how do you prevent it running up unreasonably large bills?

There was general agreement on the need for information brokering technology: an information brokers provides interesting information to its clients. There are various options for the broker to generate revenue: the clients could pay the broker; the broker could receive revenue from its information sources; the broker could receive revenue from both sources and clients.

1.9 Usability

Users are extremely important. They should not be aware of the underlying technology.

Users do not really want to be concerned with competing information providers. However, the requirement for market evolution inevitably means diversity and competition. For example, they would rather not have to lookup all airline schedules to figure out which suits them best; this would be better done automatically. Usability of information is a key issue.

The problem with many Web pages is that they have a very poor style and structure; this makes them less usable. Branding, and having a common look and feel to their Web pages, is important to Commercial organisations. There

was some discussion about style guidelines for doing this. The view was expressed that it would be difficult to legislate or impose good style and that perhaps the best way forwards was to promote good style by example. There is a tension between strict style guidelines and innovation and evolution.

Maintenance is also important: Web pages need to be maintained if they are to remain useful.

Sometimes there is a tension between usability and marketing. You might want to have very striking colours for a demonstration on a sales stand, but this might not be appropriate for everyday continual use.

1.10 Security

There was much discussion about secure HTTP and whether potential problems surrounded its use of public key encryption for authentication. There seemed to be a consensus that users need to be educated to use security measures properly. If public keys are issued and users do not protect their binding to the entity which it is supposed to identify, then it is possible for them to accidentally subvert their security. The view was expressed that electronic trade will be inherently more secure than non-electronic methods.

HTTP is stateless: the server maintains no state about its clients. This causes problems for security which requires session oriented state related to the encryption keys to be used. It is not clear what the precise trade-offs are between stateless and stateful protocols.

The issue of who issues keys and who is responsible for certification was discussed. It will not be acceptable to have one central authority.

There was some scepticism about the effectiveness of public key certification, although there was some support expressed too.

Another problem discussed was the problem of controlling copies of information. One view expressed that you cannot control copying, but you could try to make it impossible to access and use a service without paying the appropriate charge. The point was made that we may be able to learn from the music industry, since they have already had to face many of these problems.

1.11 Legal and accounting issues

Lawyers, politicians and accountants will inevitably be involve in electronic business activities. Lawyers and politicians to legislate and to resolve disputes; accountants to measure and control it. This means that electronic business needs to be explicable using paradigms and metaphors which are already familiar to such people.

The view was expressed that if accountants could not measure and understand it they would seek to outlaw it from the businesses which they control. If it costs nothing they will not worry. However, if costs something they need to be able to measure it in terms which they understand.

Prestel was cited as an early example of a commercial information service from which lessons can be learned. It eventually failed for several reasons including regulations. BT was not allowed to use the charge of the telephone call to cross-subsidise the data and provide it free.

1.12 Customer types

At least two classes of customers were identified:

- The general public
- Businesses and organisations

The distinction between customers and their agents needs to be borne in mind.

1.13 CommerceNet

CommerceNet: “an open, Internet-based infrastructure for electronic commerce, created by a coalition of Silicon Valley organizations” excited a lot of interest. Its URL is: <http://www.commerce.net/>

1.14 The World Wide Web Consortium

The World Wide Web Consortium is being established with public funding: the US Government and the European Union. Its function will be to sort out standards for the Web. Two sites will be established one in the USA and one in Europe.

1.15 What next?

The next step is to analyse the results of this workshop to see which of the issues identified will benefit from the ANSA approach.

1.16 Papers and slides presented

Copies of the following documents were distributed at the workshop:

1. ANSA WWW workshop - Introduction, Ed Oskiewicz (slides)
2. ANSA and the Electronic Market Place, Nigel Edwards (slides)
3. BT meets WWW - a clash of cultures? Dave Linton (notes)
4. Remuneration in an ODP Environment, Gray Girling (slides) (Supporting paper APM1153 is also available)
5. Global INFOWARE: The Net and You, J.D.Foss, B.C.M.Atkin (paper)
6. Brokering Service, John Moulton (paper)

Subsequently we hope to distribute further contributions from other attendees via email. Please contact Nigel Edwards (nje@ansa.co.uk) if you require these and you are not on the attendance list.

1.17 Attendees

Dave Linton (BT Labs) davel@hfnet.bt.co.uk

Gray Girling (APM/ANSA) cgg@ansa.co.uk

Alan Poppitt (BT Labs) apoppitt@jungle.bt.co.uk

Wayne Alston (ICL) wja@oasis.icl.co.uk

David Iggulden (APM/ANSA) di@ansa.co.uk
J.D. Foss (GPT) foss_jd@ncp.gpt.co.uk
John Moulton (GPT) johnm@ncp.gpt.co.uk
Yigal Hoffner (APM/ANSA) yh@ansa.co.uk
Nicola Howarth (APM/ANSA) njh@ansa.co.uk
Andrew Watson (APM/ANSA) ajw@ansa.co.uk
Dave Otway (APM/ANSA) djo@ansa.co.uk
Tim Hart (BT Labs) hart_t_p@bt-web.bt.co.uk
Dave Raggett (HP Labs) dsr@hplb.hpl.hp.com
Ed Oskiewicz (BT/ANSA) eo@ansa.co.uk
Andrew Herbert (APM/ANSA) ajh@ansa.co.uk
Owen Rees (APM/ANSA) rtor@ansa.co.uk
Nigel Edwards (HP Labs/ANSA) nje@ansa.co.uk

1.18 Acknowledgement

The author would like to thank all attendees of the workshop for their participation. Additional thanks is due to Gray Girling, Andrew Herbert, Ed Oskiewicz and Owen Rees for their help in preparing this report

