



**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

ANSA & Commercial Information Services

Nigel Edwards, Ed Oskiewicz, Owen Rees

Abstract

This note summarizes research topics identified during the recent ANSA workshop on Commercial Information Services in the World Wide Web [EDWARDS 94] which could benefit from having an ANSA approach applied to their solution.

APM.1233.00.03

Draft

25 April 1995

Request for Comments (confidential to ANSA consortium for 2 years)

Distribution:

Supersedes:

Superseded by:

1 ANSA & Commercial Information Services

This note summarizes research topics identified during the recent ANSA workshop on Commercial Information Services in the World Wide Web [EDWARDS 94] which could benefit from having an ANSA approach applied to their solution.

1.1 Agents and Brokers

These were identified as a big issue for the future as a solution to finding what's out there. A meta problem is how to find the services which can tell you what is out there (analogous to finding the public library, the bookshop or the yellow pages). A possible foundation for this is to have hierarchical sets of agents.

A problem identified for the future is that when there are many agents roaming the network, the pressure for efficiency will pressure them to share resources: a group could watch AT&T's share price together, parting company only when their instructions to buy and sell differ. (Recent Economist article [ECONOMIST 94], not the workshop.) The group itself may roam the network (analogous to car-sharing on the information highway).

A useful and necessary foundation to this is in being able to name sets of things (see Naming, §1.2). Again there is a problem of how, in the first place, to find the group watching the share price.

1.1.1 Recommendation

There is no work going on in this area at the moment. In the long term, we should consider building up a strategic core competence in the area. Consider resourcing a some (small) activity now.

1.2 Naming

Being able to name potentially changing collections of things, use of attributes to name things and naming roles were identified as important.

Naming may need to take account of location and mobility: for example a customer may be resident in Cambridge, but if he or she is visiting Chicago, a local Chicago-based service may be the most appropriate (e.g. ordering a Pizza on the network).

1.2.1 Recommendation

We already have a track record in this area with AR3 [LINDEN 93]. In AR2 [OSKIEWICZ 93] we talk about interface groups: naming potentially changing collections of things. The Remote Data Base Query Model shows how to manipulate sets of objects contained in a database. We could fold this into

some of the service manager work currently being done, which looks at maintaining membership. This means making sure this work addresses it. We probably do not need to start a separate activity specifically focused on naming, although service manager work may need more resourcing.

1.3 Caching

Identified as a key enabling technology both from the point of view of resource control and availability. There is scope for work on predictive caching strategies using AI or web trails. There are lots of consistency issues. Any solution needs to be scalable, according to Dave Raggett the big problem with current solutions is that they do not scale.

1.3.1 Recommendation

We have already done some work on replication with interface groups. We need to look at how this work and the work fits in with caching. In addition the quorum, consensus, replication and serial dependency work in the current plan is relevant here.

An activity should be started looking at this issue. It needs to be aligned with other work in the area: interface groups and service managers which deal with naming collections of things. We should start by looking at the problems with current technology by pulling over the proxy server from CERN and seeing how its caching mechanisms works.

1.4 Remuneration

This generated a lot of debate at the workshop. What is the right paradigm: electronic money, EFT, bartering? How do you make sure it is secure and not open to fraud? Billing services are needed: one stop billing, lightweight billing, auditing. How does payment granularity effect cost? It is clearly a hot button for BT, GPT and possibly Bellcore.

1.4.1 Recommendation

This has been proceeding as a background activity. We recommend that it continues to be resourced at a level where it is kept alive. EIT may be interested in this work (they are being sent the slides and paper). If they come back to us, we may need more resources for this work. Some of the work on auditing etc. that we have been doing could be focused at auditing and detecting fraud for electronic payment (we might also be able to make use of some of the team's security expertise here.)

We could exploit our location in Cambridge to build links with Ross Anderson and use his expertise.

1.5 Resource Control

This is a big issue for Web systems (see also caching above). Web trails etc. may be of use for predictive resource allocation strategies. There are two complementary approaches one is predicting and delivering the right capacity channel, the other is using (predictive) caching to minimise the demand for capacity.

There are also problems associated with (e.g. radio) stations broadcasting in web systems. You would like to avoid each user listening opening their own channel right back to the source.

1.5.1 Recommendation

We are not sure if this is being addressed by the present work. It is probably something we should be looking at. It suggests that the performance work should look at issues of higher-level orchestration of multiple entities to deliver the appropriate resources across a network.

1.6 Recovery mechanisms, auditing etc.

This was a recurrent theme (and not just in the guise of remuneration mechanisms). Particularly important was the idea of a service to resolve disputes/failures or guarantee business transactions (occurred repeatedly in the discussions about electronic money, trade associations and general fault management).

1.6.1 Recommendations

This is (at least partially) being addressed by the current work (e.g. DIPS has a notion of an audit service [OSKIEWICZ 94]). This work is probably best developed by exploring examples. A reasonable objective might be a report after exploring examples and trying to draw out some architectural principles. A possible direction for future work on caches would be to look at using redundancy to detect and resolve inconsistency. We probably need to make sure this is picked up somewhere explicitly, possibly by setting it as an objective for the MEDIC/Service management work.

1.7 Security and protection of commercial interests

A big issue which generated much discussion.

1.7.1 Recommendation

The main problem for us is that EIT have made it their own area and it is viewed by some as a “done deal”. This is particularly the case for two-party interaction. Very careful thinking as to what our contribution would be is required before we do any work on this.

Possible areas for which could benefit from the ANSA approach are: protection of multi-party interactions and a client “key-safe”. The latter is concerned with how does a client protect its key if it is read/writable or vulnerable to attack from the network.

1.8 Contracts & Federation

Contracts are important to resolve disputes and also to provide the information needed to achieve interworking when systems use different strategies (e.g. different payment strategies).

1.8.1 Recommendation

There are possible links to the work on the Advanced Trader over Orbix. It is difficult to see what can be achieved without tackling specific application examples, and then trying to abstract the general principles. We could use some of the problems discussed in the workshop to feed into the advanced trader work.

1.9 Links to existing plan

This section discusses the links between the work discussed above and the work items identified in [HERBERT 94]

1.9.1 Advance Trader

See Contracts and Federation (§1.8).

1.9.2 Work on performance

Streams, QoS Engineering, QoS Architecture, Binding Architecture, Binding Prototype are not mentioned explicitly, but nobody disputed that they were not fundamental enabling technologies. They are needed to deliver the richer kinds of media and to make it easier to engineer the fault management services (see NJE's slides presented in the workshop).

1.9.3 Medic

This is the service management work, it seems to cut across many of the issues discussed above.

1.9.4 Accounting

Clear links to the remuneration issues.

1.9.5 Cache

Obvious direct link (see also quorum, consensus, replication and serial dependency stuff which we planned to work on).

1.9.6 Interceptor & CORBA Wrapper for HTML Objects

This is so that we can do our experiments in a CORBA world. Nobody disputed that doing experiments in the CORBA world was not sensible, but nobody jumped up and down and said this is neat! But it was at the beginning of the day and may be folk were feeling nervous about voicing opinions! Another possibility is that most people were in the mode of thinking about client access rather than service access (CORBA is really about fast service provision).

There was much discussion about use of scripting to deliver services in the World Wide Web (scripts executing one or more remote services to deliver a potentially more complex service), but not about the use of CORBA technology. Was it a question of inadequate explanation of our position or is it that CORBA is irrelevant and not useful?

Note: This is not just about CORBA, it would be relevant to any remote object interaction model. Also having RPC is not enough, threads and concurrency are needed.

References

[ECONOMIST 94]

“Have data, will travel”, The Economist , May 14th 1994, pp123-124

[EDWARDS 94]

Edwards, N.J., “Commercial Information Services in the World Wide Web”, APM1220, APM Ltd., , Cambridge U.K., May 1994.

[HERBERT 94]

Herbert, A., van der Linden, R., “Focus for 1994 - 1995 activities in ANSA”, APM.1204, May 1994.

[LINDEN 93]

van der Linden, R.J., “The ANSA Naming Model”, AR003.01, APM Ltd, Cambridge, UK., February 1993.

[OSKIEWICZ 94]

Oskiewicz, E.P., Edwards, N.J., “DIPS — A Distributed Information Publishing System”, APM.1171, APM Ltd., Cambridge, UK, April 1994.

[OSKIEWICZ 93]

Oskiewicz, E., Edwards, N., “A Model for Interface Groups”, APM Ltd., Cambridge, UK., February 1993.

[THOMAS 94]

Thomas, G., van der Linden, R., “Remote Database Queries in Object-Oriented Distributed Systems”, APM.1138, APM Ltd., Cambridge, U.K., February 24 1994.

