Integrating TINA into an Internet-based Services Market

David Lewis, Thanassis Tiropanis

Department of Computer Science, University College London
D.Lewis@cs.ucl.ac.uk, T.Tiropanis@cs.ucl.ac.uk

Abstract. The deregulation of the telecommunication industry is providing the basis for a competitive, open market in telecommunications services. The TINA Consortium is developing a software architecture that aims to support both communication and information services in an open services market. This paper examines the suitability of TINA for use in an open services market based on Internet technology. An implementation of a composite, IP-based service that uses features of the TINA Service Architecture is presented, and the implications for the wider integration of TINA and the Internet in the delivery of open services are discussed.

1. Introduction

In the context of global communications, an open services market is one in which organisations, i.e., service providers, may freely sell communication-based services. Traditionally, in the telecommunications industry services have been offered to customers in a form closely related to the means of transmitting the user's signal, e.g., voice telephony, ISDN, SMDS, ATM (Constant or Variable Bit Rate services, etc.). Such service structures and their related charges are closely linked (often by regulation) to the recovery of both the costs of the initial network investment and of the operation and management of the service. Increasingly, however, services are sets of distributed software features that are presented to users as a useful whole for which they are willing to pay, e.g., IN-based services.

The Internet offers a different model of a service. The services used by the user are computer applications operating over a best effort datagram service, which is typically charged at a flat rate related to access speed. Initially, Internet users only perceived simple services such as e-mail, dns, ftp or telnet though these were rarely charged for individually by Internet Service Providers (ISPs). The advent of the World Wide Web (WWW), however, changed this picture completely. It suddenly became very easy for more or less anyone to operate a WWW server and set themselves up as a provider of services over the Internet. These services provide access to; existing services (e.g., ordering a Pizza or booking a flight), services based purely on electronic information (e.g., search engines or weather reports) or, increasingly, on communications services (e.g., Internet Phone, Video-on-Demand).

Clearly, making the mechanisms for service provision available to the widest possible business community, rather than just telelcom operators, has helped trigger an unprecedented growth in on-line services. Also, the ability of WWW-based service providers to charge for their services independently of the ISP charges has facilitated innovation and growth of these services. Access for the potential providers of such services to an easy-to-use service delivery mechanism is therefore key to any successful market in telecommunication-based services.

The Internet still suffers, however, from unpredictable performance degradation both in terms of quality of service and of service availability. These are issues that have already received much attention in the telecommunications industry. Large investments have been made in the systems that monitor the performance and manage the operation of the large networks that provide services to customers. This ensures the availability and quality of service (QoS) experienced by service customers. This has been accompanied by the development of systems gathering the data needed to charge for service usage. Typically, these systems have been specific to the network technology used and to the service offered, thus making subsequent service and management system integration difficult. In response to these problems the Telecommunication Information Networking Architecture (TINA) Consortium has been developing a model for integrated operation and management of a wide range of services (both communication- and information-based) in a network independent manner. This architecture includes a Business Model that captures the key business areas, based on the properties of TINA and an analysis of the telecommunications market. The relationships between business roles in this model provide the basis for the definition of reference points between enterprises in any TINA service implementation.

This paper gives an overview of how TINA Business Model can be applied to open service market situations and how TINA differs architecturally and technologically from the Internet. Some of these differences, and possible convergence strategies, are presented through an overview of a trial system that integrates TINA Service Architecture features and Internet-based services. The evolution of TINA-Internet convergence towards an open service market is then discussed before conclusions are drawn.

2. The TINA Business Model in an Open Services Market

The TINA business model [1] defines the following business roles.

- The Consumer role, which takes advantage of services provided by a TINA system but without the intention of providing TINA services to other business roles.
- The Broker role, which enables other business roles to locate service providers.
- The Retailer role, which is concerned with providing services to the Consumer role.
- The Third Party Provider role, which is concerned with providing services to Retailers or other Third Party Providers, but not directly to Consumers.

- The Connectivity Provider, which owns a network and provides connectivity services over it to other business roles.

A set of business relationships (see Fig 1) is specified between these business roles. TINA reference points are defined in relation to the business relationships they support. All the non-Consumer roles have self referential business relationship supporting the federation of and business co-operation between the roles.

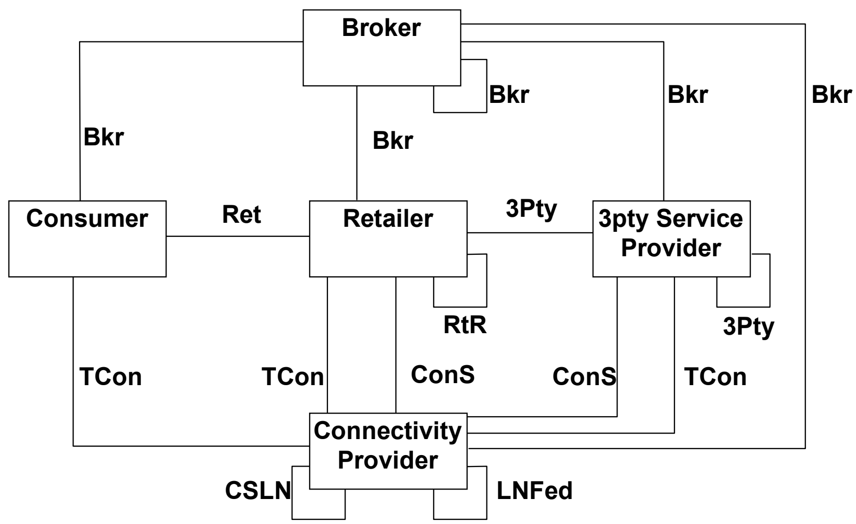

Fig. 1. TINA Business Roles and Relationships

In a business situation where a TINA conformant system is required to support inter-domain interactions, the different business administrative domains involved may be characterised by the TINA business roles they play with respect to each other. This determines the TINA business relationships they have with respect to each other. The identification of business relationships then allows inter-domain conformance specifications to be defined by the amalgamation of the reference points related to the business relationships, plus any service specific interactions required. These reference points are divided into segments, with common segments used to cover the core parts of TINA system functionality. The primary segmentation is between access functionality and usage functionality. The access segment is concerned with authentication and authorisation of users, the selection of services and the setting up the context for the use and management of the service. The usage segment is subdivided into primary usage segments covering the functionality that is the main objective of the service, and ancillary usage segments that address administrative and management functionality. This segmentation of reference point definitions enables any inter-domain reference point to be defined with the minimum set of functionality needed for the business relationships being analysed. The TINA Service Architecture [2] defines the details of interactions between components in terms of feature sets. These feature sets specify levels of functionality in a TINA system, such as basic session control or multiparty session control. Different feature sets are supported by interfaces provided by the individual software components defined in the TINA Service Architecture. These components are contained in a software architecture that supports their integration in different configurations to implement different groups of feature sets. Feature sets can be mapped to reference point segments, so that the segments for any particular inter-domain reference point should determine, via the mapping to feature sets, the minimum set of software components that would be required to conform to this reference point. The same principle can be applied between systems within a business administrative domain that are playing different TINA business roles, these are termed intra-domain reference points. This mechanism therefore provides a ready mapping from relationships in an abstract business model to the set of TINA software components needed to implement systems that provide the required TINA conformant interface functionality. In addition, TINA allows for the negotiation of the feature sets to be used between two parties at run-time. Currently, however, there is no publicly available mapping of initial TINA reference point segments to feature set definitions. When available, this will only cover the service independent session control and management functions specified in the Service Architecture. The intention is to expand this collection of mappings with further service specific feature set as they are defined by designers of specific TINA services.

3. TINA and Internet-Related Network Architectures

As well as the Business Model and Service Architecture, TINA also provides integrated Network, Management and Computing Architectures. The Computing Architecture represents one of the more novel features of TINA, in that it assumes that all systems supporting the operation and management of networks and services operate over an object-oriented Distributed Processing Environment (DPE) [3]. Most current TINA implementations have adopted CORBA as the DPE. However, most other standardisation efforts related to network technologies concentrate on control mechanisms that are protocol-based rather than DPE-based. The difference in these approaches presents a major challenge to the wider adoption of the TINA Network Architecture. The primary protocol-based technologies attracting the most attention as the basis of the future, global "Information Superhighway" is the Internet.

The Internet is based on the Internet Protocol (IP) which can operate over a wide range of underlying networking technologies, including ATM. The Internet’s proven scaleability and exponential growth point to IP as being central to any future, global communications network. As an inter-networking protocol, it is currently able to support the widest range of applications over the widest range of networking technologies. Where the Internet is currently lacking, however, is in support for communication streams with guaranteed QoS. Several solutions to this problem are currently being pursued, these being principally RSVP [4] which supports the reservation of resources for an IP stream in routers, and schemes that exploit the QoS capabilities of an underlying network such as ATM, principally using different form of label switching. These solutions are again the subject of intense interest from equipment manufacturers and ongoing activity in the relevant standards bodies, i.e., the Internet Engineering Task Force and the ATM Forum.

We can observe, therefore, that if TINA is to be widely adopted in the near to medium term it must be able to accommodate current and evolving protocol based communication control mechanisms, and in particular QoS-aware IP networks. TINA already assumes some use of IP, through its adoption of CORBA and its Internet Inter-Orb Protocol for its DPE. CORBA, however, does not currently address the transport of multimedia streams with guarantees QoS. Instead TINA addresses this in a technology independent way, through the abstractions of the Network Resource Information Model. Control of TINA services relies on the following session separations:

- An access session, which controls user's ability to access services offered by providers;
- A service session, which controls the communication between groups of users and the configuration of multimedia data streams between them and relevant service components and
- A communication session, which allows centralised service session components to specify communication links between logical network nodes involved in a service session.

The communication session therefore offers Service Architecture components a single point at which connections over several network domains can be requested as logical connection graphs. However, in keeping with it telecommunications orientation, this network model is still fundamentally connection-oriented, making it cumbersome to match this model well to the best-effort, connectionless datagram model of IP protocols. The introduction of flow identifiers in IPv6, and the use of soft connection state in RSVP, could be argued to make IP network appear more "connection-oriented" and therefore closer to the connection-oriented QoS-aware TINA network model. One major difference remains however, and that is in the multicast capabilities of IP networks. This capability offers direct support for user initiated participation in multiway communications, and is the basis of much of the interactive multimedia communication currently being performed on the Internet [5]. Multicast is implicitly supported by RSVP, and is a key feature of other IP QoS guarantee schemes.

Connection oriented networks typically provide only point-to-point or point-to-multipoint connection primitives, and thus tend to push the control of multiway group communication to centralised point at a higher architectural level. This can be seen in the TINA Service Architecture, which is heavily influenced by the need to control groups of connections involved in multiway communications. The requirement for centralised control of group communication can be relaxed in a service architecture where support for multiway group communication is provided by IP multicast. Therefore if a TINA system was to be deployed over one of today's IP network, e.g., a corporate intranet, the lack of support for QoS and the implicit support for multicast may enable the service architecture to be simplified, with control of group communications devolved to IP multicast functionality. This was the approach adopted in the Prospect project [6] and is described further in the next section.

One problem now recognised by the Internet community is the need for networks supporting guaranteed QoS to implement some negative incentive to users from hogging resources. In a global Internet operated by commercial providers, charging for different levels of QoS is seen as a way of providing this negative incentive, as well as funding the expansion of the network and, hopefully, making a profit. The Internet, however, currently lacks a common agreed mechanism for charging for different QoS requests and authenticating the requesting entity so that they may be billed for its provision. The TINA Service Architecture provides a generic model for offering, accessing and charging for different services with different profiles that may specify permitted ranges of QoS. The interaction of the access session with subscription management ensures only authenticated users can access specific services. The interaction of the service and communication session with accounting management provides a generic mechanism for collecting usage metering data and combining it with tariffing information from subscription management in the generation of billing information. The design of the Service Architecture, therefore, embodies much of the experience of the telecommunications industry in metering and charging for different services at different QoS levels, and could therefore prove useful in providing QoS guaranteed IP services..

4. Example: A Composite Tele-Education Service

Here we discuss an open service business scenario that was implemented in the Prospect project. The scenario consisted of an educational service provider that wished to start delivering its product to customers remotely. This organisation had the skills needed to develop and deliver its educational material through interactive WWW sites, but did not have the knowledge or skills needed to develop, operate, manage and charge for the communications services needed to deliver its product to the customers. The provider therefore used the services of two types of providers in delivering a composite, IP application-based tele-education service to its customers. The first type, termed a Multimedia Tele-Service provider, provided IP-based services using multicast conferencing and WWW applications. The second type, termed a Virtual Private Network provider, supplied a managed inter-domain intranet (sometime called an extranet) connecting the server sites of the TES and MMTS providers with those of the TES customers. In terms of the TINA Business Model the following business relationships were present (as summarised in fig 2):

- A Ret relationship between the TES customer in the Consumer role the TES Provider in the Retailer Role.
- A 3pty relationship between the TES provider in the Retailer role and the MMTS providers in the Third Party Providers roles.
- A ConS relationship between the TES provider in the Retailer role and the VPN Provider in the Connectivity Provider role.
- A TCon relationship between the VPN Provider in the Connectivity Provider role and the TES Provider, the MMTS Providers and the TES Customers as nodes on the physical connection graph.

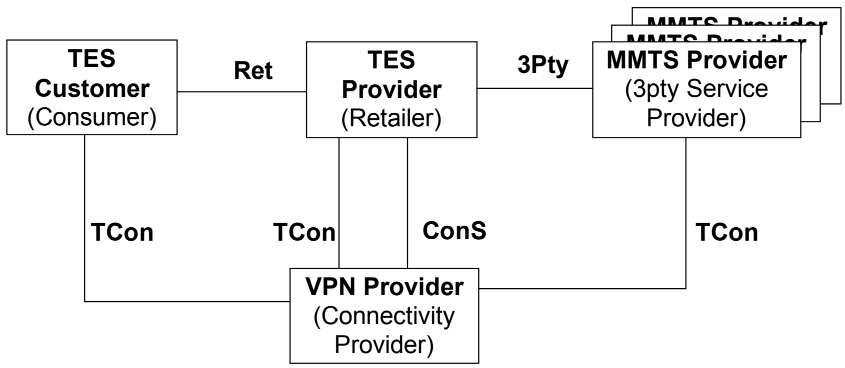

Fig. 2. TINA Business Roles and Relationships in Prospect

The design of the stakeholder systems was not intended to conform to the specifications in [2], and was largely influenced by an earlier release of this document. In the Prospect model the VPN Provider provided a conventional IP datagram service, supporting IP multicast but not the transport of multimedia data streams with guaranteed QoS. There was therefore no direct interaction between service session components and the VPN service via the ConS reference point. The VPN provided network configuration services, but these consisted of management activities to add or remove customer premises networks or to alter the bandwidth available on the IP backbone (the latter being implemented using an ATM Virtual Path management service). As end user applications were able to request the network directly to join or leave multicast groups, the need for a centralised group communication control was removed. Therefore, there was no need for a direct link between the service session and the communication session and no need for any centralised control component in the service session to perform this link or to manage group communication. This resulted in the Service Session Manager being removed for this implementation of the TINA Service Architecture. Instead only the User Service Session Manager (USM) component was used in the control and accounting of individual users' usage of services. This component was implemented differently depending on the service being provided. For Web-based services the USM was implemented as a http proxy connecting the user to a Web server. For multicast-based conferencing applications the USM provided the mechanism for distributing multicast group rendezvous information and session-specific encryption keys. Subscription and accounting management, therefore, still play a key role in managing the services at the access and service session level, regardless of the absence of network QoS guarantees.

5. Discussion

The above example revealed two important points related to the TINA Business Model. Firstly, no Broker role was present. For users, access to TES and MMTS services was performed via CORBA capable Java applets running in a WWW browser. Services were located by links to a WWW-page containing an applet that operated as a client to a User Agent in the relevant provider domain. It is assumed that this page is located by existing means, i.e., a conventionally advertised URL or via a WWW search service. References to specific provider interfaces were located in a globally available CORBA Naming Service keyed from information in the providers' WWW-page. As these existing location services proved adequate and familiar to the user, a specific TINA conformant Broker role was therefore not regarded as necessary.

Secondly, the TES provider offered no primary usage functionality since the TES home-page and Web-based educational material was hosted by one of the MMTS providers. Hence both the primary service usage and the access to this usage were delegated by the TES provider to the MMTS providers. This was possible due to the propagation, by the TES Provider, of certain customer subscription management activities to the MMTS business administrative domains. For example, when a user was authorised or barred from a TES Provider service, the same operation was performed by the TES Provider for that user on the MMTS services that composed that TES service. Service usage in this scenario, however, constitutes an inter-domain reference point between a Consumer role and a Third Party Provider role that is not directly mapped onto a business relationship in the TINA business model. This matches well to the business situation being considered in Prospect, where the TES Provider, in the Retailer role simply wants to directly manage the business relationship with the TES customer, but wishes to leave the control of service delivery to more specialised providers of communication and information based services. This implies that the delegation of service control to inter-domain interfaces that are not covered by a direct business relationship, e.g., between the Consumer and the Third Party Provider, may be of more significance in an open services market, than currently indicated by the TINA Business Model.

Considering possible migration strategies between TINA and the Internet, the integration of QoS guarantee mechanisms and call models appear promising in the near term. As discussed above, it is likely that guaranteed QoS provision in the Internet will be delivered through user application signalling rather than the centralised network control model adopted in TINA. The example of the previous section demonstrates how the TINA service architecture is useful for providing authenticated access to and accounting of IP-based service session at the application level. The principles of delegating network control to the user application used here should also apply to the use of IP-based QoS requests. However, once QoS guarantees are provided by an IP service then network level access control and accounting is also required. This would therefore require some interaction between the IP QoS request call admission control function and TINA components in the Connectivity Provider role, possibly in line with the work of the RSVP Admission Policy working group in the IETF [7]. The call admission control function would need to check with the subscription management component to check for any policy applied to this service usage. The call admission control component could then record successful QoS request and the duration of the flow using the accounting component. Though the request comes from the user application, the authorisation token could come from another entity, e.g., the Retailer, so that the charges for use of the network could be sent to parties other than the user.

TINA employs a call-model in the definition of the Service Architecture. This call-model requires a number of TINA specific components on both the retailer's and the consumer's domain. Some of these components have to be on the end-user's terminal. On the other hand, the call-model of the Internet world employs protocol-based exchanges for session invitations and for session announcements. The session announcement protocol (SAP) [8] can be used to multicast session announcements to a number of users who are members of a specific multicast group. The session initiation protocol (SIP) [9] provides for directing session invitations to individual users.

An important step towards an integration between TINA and the Internet is to allow current Internet users to announce and initiate sessions with the existing Internet protocols in a TINA environment, without having to introduce TINA specific components on the end-users' terminals. This requires a mapping between Internet invitations and TINA invitations. This issue is being investigated in the ACTS project VITAL.

6. Conclusion

The telecommunications market is a rapidly changing one, so any attempt to develop a model for telecommunications systems based on a fixed business model and specific technology assumptions runs the danger of producing architectures that are inflexible and therefore short-lived. Attempting to apply the TINA to an Internet-based open service example yielded several interesting insights into the integration of TINA and the Internet.

- The Internet already provides a wide range of service location facilities through advertising and the Web. It is not clear therefore that a specific Broker role will be necessary.
- The delegation of service usage interactions to inter-domain reference points not covered by direct, contractual business relationships may be more significant than current indicated in the TINA Business Model.
- The centrally controlled, connection oriented network model adopted in TINA fits poorly with the Internet communication model. The design of TINA therefore needs to accommodate the delegation of control of the network from a centralised mechanism to the user application driven one.
- TINA needs to recognise the persistence of signalling for QoS requests, and a systematic way of integrating such protocol-based mechanisms with the DPE approach is required. The adoption of the web-browser as the basis of the user

application should support the integration of Java-based CORBA components with protocol specific components that is needed as part of such an approach.

It seems possible, therefore, that despite some fundamental architectural differences, parts of TINA could be applied be successfully deployed in the Internet in the near term. In this way, the use of TINA access and service session concepts and service management components can bring the advantages of open service management interfaces and reusable software components to ISPs.

7. Acknowledgments

The work presented in the paper was conducted with partial funding of the European Commission under the Prospect project (number AC052). The authors would also like to thank all members of the Prospect team for their collaboration in this work. The views expressed here are not necessarily those of the Prospect consortium.

References

[1] Yates, M., Takita, W., Demoudem, L., Jansson, R., Mulder, H., TINA Business Model and Reference Points, v4.0, TINA-C, May 1997

[2] Abarca, C., et al., Service Architecture, v5.0, TINA-C, June 1997

[3] Graubmann, P., Mercouroff, N., Engineering modeling Concepts (DPE Architecture), TINA Baseline Document TB_NS.005_2.0_94, December 1994

[4] Zhang, L., Deering, S., Estrin, D., Shenker, S., Zappala, D., RSVP, a new resource ReSerVation protocol, IEEE Network, vol. 7, no. 5, September/October 1993.

[5] Handley, H., Crowcroft, J., Bormann, C., Ott, J., The Internet Multimedia Conferencing Architecture, Internet Draft, draft-ietf-mmusic-confarch-00.txt, July 1997

[6] Lewis, D., Tiropanis, T., McEwan, A., Redmond, C., Wade, V., Bracht, R., Inter-domain integration of services and service management, Proceedings of the 4th international conference on intelligence in services and networks, IS&N'97, Cernobbio, Italy, May, 1997

[7] Yavatkar, R., Pendarakis, D., Guerin, R, A Framework for Policy-based Admission Control, Internet draft, draft-ietf-rap-framework-00.txt, November. 1997

[8] M. Handley, SAP: Session Announcement Protocol, Internet-draft draft-ietf-mmusic-sap-01.txt, November 1996.

[9] M. Handley, H. Schulzrinne, E. Schooler, SIP: Session Initiation Protocol, Internet-draft draft-ietf-mmusic-sip-04.txt, November 1997.