Infrastructures for Mobile Government Services

Mícheál Ó Foghlú
Telecommunications Software & Systems Group
Waterford Institute of Technology,
Waterford, Ireland
E-mail: mofoghlu@tssg.org,
http://www.tssg.org/

Abstract : *This paper describes a range of infrastructural issues that impact on the way that mobile government services are being and will be designed and deployed. In this sense, mobile government services are just like many other areas of mobile computing, dominated by the push to deploy Internet technologies on wireless networks. The first issue that is addressed is that of the ownership of the infrastructure. It is argued that the concept of Open Access Networks (OANs) is an important one, especially for publicly funded wireless networks. The next issue addressed is the use of IPv6 instead of IPv4 as the main transport technology. Political and technical reasons for the choice of IPv6 are discussed. Finally the paper addresses the software used to develop services and touches on the debates on appropriate middleware architectures for mobile services.*

Keywords: IPv6, Mobile Internet, Open Access Networks, OAN, middleware, web services, REST

1. Introduction

Infrastructures for mobile computing in general, and mobile government as a sample domain, made up of wireless networks and mobile access devices accessing flexible software services, are emerging as a key area for research in computing. The Irish HEA-funded research programme M-Zones (M-Zones, 2005) advocates a general approaches to the exploration of the management of these infrastructures and services. With all wireless networks there is the option of deploying one's own infrastructure, and managing it, or utilising an external infrastructure, usually belonging to a telecommunications operator. If one is supporting users on a fixed site, the former is a valid option, if one is supporting users roaming more widely the latter often becomes preferable.

This paper addresses a series of underlying infrastructural issues for mobile services, and for potential generic management structures for these services. This infrastructure is built on existing and emerging network protocols, and on higher level approaches.

The aim of this paper is to proselytise a number of key decisions that the author feels would help progress the current plethora of incompatible services towards the potential vision of a network of integrated managed zones. This draws on research in the areas of wireless networking, data networks, telecommunications networks, and on research into applications and management services over these networks.

2. Open Access Networks

Although it may be perceived as being slightly peripheral to the central theme of infrastructure issues, the concept of OAN (Open Access Networks) is an interesting one with a very specific message for wireless networking research, particularly where there is the possibility that public funding (local, regional or national government sponsored) is used to deploy the networks for the use of a wide

audience. The vision of Open Access Networks is to conceive of a world where the networks are owned/operated by neutral entities (potentially subsidised by regional or national governments in peripheral regions, potentially paid for by a small levy on the service providers in more populated areas) who allow any number of operators to use the infrastructure to offer competing or complementary services. The reduction in the cost through the avoidance of the deployment in many separate infrastructures can lead to more cost effective solutions for end users.

One pioneering institution in the area of OANs is KTH (The Technical University of Stockholm, Sweden). KTH has deployed a metropolitan WAN based on OAN principles. The network comprises both wireless access points (currently 802.11b) and broadband links (one partner in the enterprise is the Swedish housing authority who own a large percentage of rental accommodation in Stockholm, and who have deployed the OAN into newly refurbished buildings). In a recent conference paper (Battiti et. al. 2003) the authors describe the concepts behind this concept, and recount the practicalities of the deployment in Stockholm.

The concept from the end users' perspective is that they can use a single infrastructure (whether wired or wireless) to connect to their operator of choice. They can even switch from one operator to another, currently using a web-based interface. The underlying principle is that the system uses a DHCP proxy system; it works with IPv4 and IPv6, to redirect any particular device's DHCP request to their operator's DHCP server. Then, local firewall rules prevent the allocated IP address form any access until the end user is authenticated. Authentication is delegated to the operator, so an operator could allow relatively lightweight authentication, or insist on more heavyweight authentication (e.g. RADIUS, DIAMETER, or smart-card based).

The TSSG have deployed a version of the KTH software (available as an open source download from http://www.StockholmOpen.net) to allow the potential for multiple operators within the TSSG wireless network. Currently only a single such operator exists (the TSSG itself) and the research group has not established business relationships with ISPs and WISPs as has been done for the StockholmOpen network. The potential is to link this infrastructure to the emerging metropolitan WANs in Ireland (funded by the Department of Communications & the Marine) for example the SERPENT project in the South East Regional Authority (Ó Foghlú, 2002). This initiative was championed by the South East Information Society Strategy (SEISS), and has now led to the established of a large number of Metropolitan Area Networks (MANs) with an open access policy. The network management of all of these MANs has been subcontracted to a company called E-Net (E-Net, 2005), effectively a pan-Irish experiment in open access networks with central government and local government funding. The potential is that mobile solutions can be built on similar principles using the metropolitan infrastructure as an optical backbone.

The alternative to the OAN approach is to continue with existing models. This means, in the telecommunications-linked networks of 2.5G and 3G, working with existing operators. It means, in an Internet/computer network scenario, relying on a local company or a specific ISP or WISP with the overhead of separate competing infrastructure, probably not offering any service migration between them, within the same location. In general, this paper suggests that mGovernment projects could benefit from an analysis of the OAN debate, and consideration of the provision of a publicly supported infrastructure as part of the project. Public sector network investments are ideally situated to evaluate the benefits of OANs and to provide these facilities in locations where commercial alternatives may not be viable.

3. IPv6 and the End-to-End Argument

At its simplest the case for IPv6 is based on the available range of addresses for Internet devices in IPv4 (32bit) versus IPv6 (128bit). It is true that, using a NAT-enabled infrastructure, IPv4 can continue to be deployed in Western Europe and the developed world for the near term. Here there are enough IPv4 addresses for ISPs to assign limited addresses to the publicly accessible machines and hide the complexity of the real diversity of internal machines (as they are given a private Class A address range such as 10.x.y.z). However, it is already the case that IPv6 has become the Internet protocol of choice for Japan, Korea, China, India, and the Far East in general (and to a lesser extent of South America and Africa as well). Population pressures, and the political residue of the differential treatment of these regions, have prompted strong governmental and industrial support for IPv6. In the Far East the IPv6 address allocations are perceived as being more culturally neutral and as essential to overcome the immediate addressing concerns today. This is because of the control the USA had over the allocation of IPv4 addresses, and the perception that effectively this allocation favoured the western economies.

The argument so far has made no mention of mobile devices. When these are introduced to the equation, timescales are shortened further, a large address range is even more important, as there is the potential for multiple devices per person (rather than just one or two, a desktop and/or laptop). Recognition of this was a major factor in the selection of IPv6 as the protocol of choice by the 3GPP for the deployment of a world-wide UMTS 3G mobile telephone network. So there is no need to look further than this one simple argument: an in a world wide mobile network, based on IP protocols, it makes sense to use IPv6 rather than IPv4. Any research looking to the future of mobile data services should make this choice now.

The counter-argument to IPv6 is relatively simple: what does IPv6 give that forces users and providers into a migration to IPv6? The answer for many is nothing. Whilst there are more addresses, many can survive happily with existing IPv4 address ranges, especially in Europe and North America, and use NAT to create privately addressed subnets where needed. These counter arguments are as much about the financial justification for upgrading corporate infrastructures as they are about IPv6. Meanwhile, the more populous countries where the issues that IPv6 solves are more pressing may push ahead of us in IPv6. If this prediction is correct, as it already seems to be becoming, this then means there is a business argument, for companies whish to develop products for deployment in these, by definition, large markets, the maybe they should push IPv6 more at home as well.

There are other reasons why IPv6 is the obvious choice for mobile networks. IPv4 requires a considerable infrastructure to allow the assignment of addresses on the network: a router and a DHCP server are needed. In IPv6 this auto-configuration is part of the basic infrastructure, supported by routers themselves. It is close to zero management in that once the router knows the network prefix, the addresses for devices are assigned automatically. Furthermore these are predictable as they are derived from the MAC address of the Ethernet device. These are called Global Addresses. Even if no router is available, the device configures itself with a Link Local address, and may be able to tunnel over an IPv4 network. While such addresses may be useful for local traffic, effectively for clients, these addresses are not useful as a generic global identifier for devices, thus missing out on one of the key benefits of IPv6 (the use of the address itself as a unique global identifier of an end-point). The trend has been in many areas of distributed computing for every node to be a server as well as a client (especially so in peer-2-peer networking) and this then raises the issue of uniquely identifying the server, potentially in a global context.

Throughout the history of the Internet there has been a debate about the potential benefits of end-to-end communication (Saltzer et. al. 1984), (Reed et. al. 1998), (Clark et. al. 2002). At its simplest, the

argument is that if every device on the network has its own unique address, this makes it very easy for any device to offer a service to any other device. In TCP/IP such services are identified by the combination of the IPv4 address and a port number. So if one machine wants to offer a service, it simply has to advertise its address and the port number it is offering the service on. In most senses this architectural purity has been sullied by the use of NAT (Network Address Translation) in IPv4. Ostensibly used to try and extend the range of IPv4 addresses by using private address spaces inside organisations, and then only exposing a limited number of addresses outside, this mechanism also breaks the open nature of the Internet. Now every machine cannot offer every other machine on the network a service.

One positive potential for the use of IPv6 is the restoration of an end-to-end Internet where every endpoint has a unique address, thus allowing every machine to potentially offer a services to every other machine without the "problem" (from a service access point of view) of firewalls and NAT. However, re-establishing this paradigm would raise huge security issues. Currently many companies and ISPs use NAT as a way of hiding the real identity of endpoints inside their domains (changing all outgoing packets to a single publicly visible IPv4 address). NAT is used not only to solve the issue of the shortage of IPv4 addresses, but as primitive security tool. Of course, many applications have now been engineered to tunnel over the protocols that are allowed through NAT gateways and firewalls (primarily web protocol http on port 80 of the remote server). The normal response of the IPv6 community is that more secure networking is possible with IPv6 as IPSec is mandated, thus providing the potential for a network-level securely encrypted session. Of course, this may not address the requirements of central control in companies and ISPs as to what kind of traffic is and is not allowed into and out of their networks. This debate is important and as new solutions emerge, smart space frameworks and services need to be able to integrate. There may be an IPv6 world were it is equally difficult to get through corporate choke points of various kinds. For telecommunications companies, they may try and restrict IPv6 services on 3G networks to ones offered by themselves and by paying partners, rather than opening up their 3G networks to all IPv6 Internet services (to a large extent this was the way WAP was approached by most operators). There may be commercial as well as security and address-range rationales for network choke points.

Theoretically IPv6 promises easier provisioning of QoS (Quality of Service). Currently there is no real QoS on the public Internet, though there are many deployments of QoS frameworks (most popularly MPLS-based) on internal networks, especially when used for VoIP traffic as an alternative to public telecommunications networks. mGovernment frameworks could potentially allow for the potential for IP-based QoS provision and negotiation (e.g. IntServ, DiffServ and MPLS). For the conceivable future, QoS on the public Internet is a pipe dream, and network engineers will continue to solve problems by deploying more bandwidth rather than engineering QoS differentiation, as has been done throughout the history of the Internet and of LANs. The TSSG have been a partner in the EU FP5 IST project Intermon (Intermon, 2005) addressing the issues of Interdomain quality of service, primarily with an IPv4 focus.

A very promising facility offered by IPv6 is the use of Mobile IPv6. When Mobile IPv6 is deployed, this framework can allow for the transfer of a session from one IPv6 endpoint to another relatively seamlessly. In contrast Mobile IPv4 is routed via the original home node and so can be very inefficient. Clearly it is important for this type of macro-mobility to allow users to move from one place to another (potentially administered by different authorities) and to continue to use a service without interruption. Of course the issues of coverage and of organisational relationships to allow such roaming are larger than the issue of being able to negotiate a new IPv6 address and continue with a service originally accessed from a different IPv6 address. As Internet applications have traditionally not included large elements of network management, there is much work to be done in establishing such interrelationships. It is more likely to be tackled in the 3G world (where potentially revenue can

fund such research) rather than in the open access WiFi world where there is less incentive to creating roaming agreements, and less agreement as to who the operators are and what their responsibilities are.

Currently, it seems as though there is a critical mass of IPv6 deployment, particularly in Japan and Korea, but also in the cluster of EU funded projects in the Fifth Framework (IPv6 Cluster, 2002), (IPv6 Cluster, 2003), and new projects in the Sixth Framework. This level of activity will hopefully see solutions to any outstanding barriers to the deployment of IPv6. These barriers are often to do with the commercial availability of equipment on the market to deploy in an IPv6 infrastructure. Therefore, it seems reasonable to suggest that any research into mobile IP networks should focus on the use of IPv6 as well as IPv4 (or to the exclusion of IPv4 where appropriate). Awareness of issues relating to IPv6 does not seem to be there in the eGovernment and mGovernment communities, but, just as with all IP services, the issues of end-to-end are very important in potential government applications and services.

4. Middleware Architecture

It could be argued that an equally important part of the infrastructure for mGovernment solutions is the higher level software infrastructure (the middleware framework). Here there are a larger number of alternatives vying for attention. Perhaps the lessons are best learned from the Internet itself, where simple systems built on a stateless protocol (http) proved to be the most robust and scalable. Despite not having a formal theory for describing such applications they prospered, initially through CGI-based web applications and later through more sophisticated server-side memory sharing models. An architectural description for this type of distributed application was retrofitted later: REST (Representational State Transfer) (Fielding et. al. 2002). The principle is that loosely coupled elements on a network could scale up to provide useful services.

So, it may be possible to build mGovernment using light-weight architectures like the web itself (though even an embedded web server may be too heavy-weight for some sensor devices). In contrast to this light-weight loosely-coupled approach, middleware research, at least throughout the 1990s, has been dominated by object-oriented heavy-weight engineering approaches (CORBA, J2EE, and Microsoft's COM/DCOM). To some extent it could be argued that the emerging interest in Web Services concepts (SOAP, WSDL, and UDDI) represents a compromise between these two paradigms. Despite the literal meaning of SOAP (Simple Object Access Protocol) the architecture actually defines message passing between services, rather than passing object references (as CORBA did) that then imply stateful architectures.

Therefore it is promising to envisage the construction simple lightweight RESTful mobile services over IPv6. One problem with this approach is that although IP by itself does not impose a heavy processing burden, though imposing security in the form of encryption can do so (e.g. IPSec), the use of verbose XML-based or other verbose text-based approaches to pass information over wireless networks with limited bandwidth may not be ideal. The bandwidth itself may be used inefficiently, and the nodes may not have the computing power required to parse and process these messages. The usual solution to this is to use various types of distributed processing nodes or gateways to wrap non-standard or less powerful client devices. This pragmatic approach is useful for building systems, but harder to treat orthogonally.

The use of the XML-based systems such as VoiceXML offer a very cheap way to deploy what were previously expensive IVR systems that may actually be an excellent way to deliver information to people using fixed and mobile handsets. Thus using simple Internet architectures to develop traditional voice delivered systems could be a fertile area for deployment of converged systems.

5. TSSG Experiences

The Telecommunications Software & Systems Group (TSSG) in Waterford Institute of Technology has been involved in a series of projects that relate to the emerging area of mobile services both from a telecommunications perspective (including 2.5G and 3G), and from an Internet services perspective. The group has lead the Opium (http://www.ist-opium.org) project testing the interoperability of mobile telecommunications services and the AlbatrOSS (http://www.ist-albatross.org) project establishing a software framework for developing and managing mobile services based on NGOSS. The group has done and is still participating in a series of projects relating to mobile and fixed IP-based services: the Torrent project investigated issues of ISP-home gateway communications using IPv6; the SEINIT project defines new security models and policies to address the new issues of the pervasive computing world; the Daidalos (http://www.ist-daidalos.org) project is dedicated to the design of advanced network infrastructures and access technologies for location-independent personalised communications services.

As the widest deployed mobile computing systems in the world today, mobile telecommunications systems highlight many of the key issues for mobile services (including some aspects of context awareness such as personalisation and location and also the area of management of mobile services). However, as these approaches are centrally controlled by a limited number of operators in each country, these types of system only offer only one perspective on the emerging mobile computing environment. An alternative perspective comes from the more organic growth of various types of localised wireless hot spots and also from the evolution of enterprise computing from predominantly wireline networks to a combination of wireless and wireline networks. Here approaches based on a centralised standards process with one mandated approach may find it harder to win out as has happened in the development of Internet-based services that has fostered a huge variety of approaches.

In order to communicate the results of these projects and to create a focal point for discussion and debate around the infrastructure and management issues relating to mobile services the TSSG have developed an integrated demonstrator incorporating results from many of these projects. This first version of this demonstration, sponsored by O2, was presented in the Irish eWeek (26th-30th April 2004) event in Ireland's Digital Hub (c.f. http://www.o2home.ie). Further versions of this core demonstration set are being developed with more focus on the internal management structures for services and users. The M-Zones research programme and other ongoing projects are contributing to this evolving demonstrator.

Our experience to date in Government projects has been more limited, with an e-procurement system for local government E-PROC (Interreg IIIb) and a system for development of unified identification for government services RISER (eTEN). Our experience has shown that, from a technology perspective, the same set of technologies and issues arises in various domains, and researchers in mGovernment should be open to work in other domains to maximise the cross-fertilisation of good ideas. However, all projects gain traction for links to pragmatic decisions that do relate to the specific domain, but "Government" is a large nebulous domain with many such sub-areas.

6. Conclusion

This paper has argued that wireless network infrastructures should potentially be regarded as overlay networks sharing common physical network infrastructures thus maximising use of resources. It has proposed the Open Access Network (OAN) as a potential model for such a structure, built using open source components such as those developed by StockholmOpen.net.

The paper has also argued that the case is clear for the use of IPv6 as the network protocol of choice for mobile service interoperability at the transport layer. Whilst there are still some potential issues with the use of IPv6, especially relating to security, the advantages for mobile networks with potentially large number of devices make the case for its adoption compelling.

The paper has addressed the potential reasons why the light-weight approaches used in Internet services may prove equally flexible in mobile service developments. Thus we can expect to see the use of traditional web server style applications and of the newer web services style services.

Arguably, the whole area of mobile computing is much more in flux than any other area of computing. Wireless technologies have been dubbed the next disruptive technology after the Internet, where a disruptive technology stimulates very fast development and progress through an innovative new approach challenging previous approaches. It is not yet clear what hardware and software platforms will dominate what middleware architectures for software will dominate, or even what wireless networking protocols will dominate this area. Despite this uncertainly this paper has attempted to highlight some emerging issues relating to infrastructures for these systems. Expect to see various attempts at synthesising these approaches appear in commercial products in the next few years, and provide the core infrastructural elements for mobile government services, as well as mobile services for other domains.

References

Battiti, R. Cigno, R.L. Orava, F. and Pehrson, B., (2003). Global Growth of Open Access Networks: from War Chalking and Connection Sharing to Sustainable Business WMASH'03, September 19, 2003, San Diego, California, USA.

Clark, D.D. Wroclawski, J. Sollins, K.R and Braden, R., (2002). Tussle in cyberspace: Defining tomorrow's Internet SIGCOMM'02, (August 2002) pages 19-23, Pittsburgh, Pennsylvania, USA.

E-Net, (2005) Map of Regional Metropolitan Area Networks managed by E-Net ,URI: http://www.e-net.ie/regional_rollout.htm [Last visited: 2005-02-25]

Fielding, R and Taylor, R., (2002). Principled design of the modern Web architecture *ACM Transactions on Internet Technologies* 2, 2 (2002) pages 115–150.

Intermon, (2005) *EU FP5 IST Project Website* URI: http://www.ist-intermon.org [Last visited 2004-04-29]

IPv6 Cluster, (2002). *IPv6 Research and Development in Europe* (on-line version available at URI: http://www.ist-ipv6.org/ [Last visited 2004-04-29])

IPv6 Cluster, (2003). *Moving to IPv6 in Europe* ISBN 3-00-011727-X (on-line version available at URI: http://www.ist-ipv6.org/ [Last visited 2004-04-29])

M-Zones, (2005) *M-Zones Programme Website* URI: http://www.m-zones.org [Last visited: 2005-02-25]

Ó Foghlú, M., (2002). Regional Initiatives in the South East of Ireland in *Challenges and Achievements in E-business and E-work* Edited by Brian Stanford Smith, Enrica Chiozza & Mireille Edin (Volume 1) IOS Press, Amsterdam, 2002 ISBN 1-58603-284-4

Reed, D.P. Saltzer, J.H and Clark D.D., (1998). Comment on Active Networking and End-to-End Arguments. *IEEE Network 12*, 3 pages 69–71.

Saltzer, J.H. Reed, D.P. and Clark D.D., (1984). End to End Arguments in Systems Design *ACM Transactions on Computer Systems* 2, 4 (November 1984) pages 277-288. An earlier version appeared in the Second International Conference on Distributed Computing Systems (April, 1981) pages 509–512.

Author biography:

Mícheál Ó Foghlú is the Research Director of the Telecommunications Software & Systems Group in Waterford Institute of Technology. This group, founded in 1996, has engaged in over 40 basic and applied research projects, winning nearly 20 Million EUR of funding. The group now has 56 staff and students and acts as a regional catalyst for ICT innovation in the South East. Mícheál's research interests centre on the use of next generation networks to deliver flexible services. This includes an interest in IPv6, IPv6 QoS and Security, mobile IPv6 and web-based technologies to build flexible services running over IP networks.