

Infrastructure Requirements for Smart Spaces and Managed Zones

Mícheál Ó Foghlú

Telecommunications Software and Systems Group (TSSG)

Waterford Institute of Technology (WIT)

mofoghlu@tssg.org

Abstract

This paper discusses the underlying infrastructural requirements for Smart Spaces and Managed Zones. The focus is on generic technologies rather than specific solution to specific problems. At a political ownership level, who owns the infrastructure is key. This paper addresses the issues of Open Access Networks (OANs) with pointers to the issues they raise. Such networks allow the reuse of a single underlying infrastructure for multiple purposes, allowing logically distinct smart spaces to overlay each other in a single physical wireless or fixed IP-based network. At the network protocol level, one key assumption is that smart space infrastructures are based on the Internet Protocol (IPv4, IPv6). This paper will argue that, for mobile devices, IPv6 is the obvious choice. The reasons include the available address range for a large proliferation of devices, the ability to allow p2p (peer-to-peer) interactions and lightweight services end-to-end across such IPv6 networks, the ability to allow auto-configuration (allowing flexible lightweight management), and the ability to support redirection of traffic to new locations (mobile IPv6).

[The M-Zones research programme has been funded by the Irish Higher Education Authority, PRTL Cycle 3]

1. Introduction and Context

Smart Spaces, made up of wireless networks and mobile access devices accessing flexible software services, are emerging as a key area for research in computing. Right from the early days of ubiquitous computing [Weiser 1993] until today there are a relatively large number of projects and research programmes looking at individual environments where users can interact with a specific smart space using a customised device, customised services and a customised network infrastructure. Perhaps more interesting is the smaller number of projects and research programmes investigating the emerging standardisation of this whole field. The Irish HEA-

funded research programme M-Zones¹ is one of these more general approaches.

This white paper, part of the scoping exercise of the M-Zones programme, addresses a series of underlying infrastructural issues for smart spaces, and for potential generic management structures for these smart spaces. 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 authors feels would help progress the current plethora of incompatible smart spaces towards the potential vision of a network of integrated managed zones each comprising subordinate smart spaces. 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 smart space research. In general, most smart spaces deployed in the world today deploy their own network infrastructure to support that specific smart space. In this sense they are following in the footsteps of earlier telecommunications networks where the network provider and the service provider are usually the same entity. The challenge of Open Access Networks is to conceive of a world where the networks are owned 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).

One pioneering institution in the area of OANs is KTH (The Technical University of Stockholm,

¹ <http://www.m-zones.org>

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 from 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 authentication).

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 like this type of 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, a project which was proposed by the South East Information Society Strategy (SEISS) [Ó Foghlú 2002a].

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 infrastructures, probably not offering any service migration between them, within the same location.

3. The case for IPv6 (and issues raised)

This whitepaper restricts itself to the IP-layer and above of the OSI network protocol stack, i.e. network and transport layers and above.

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.

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 now there is the potential for multiple devices per person (rather than just a one or two, a desktop and a laptop). Evidence of this can be seen 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: 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 should make this choice now.

There are other reasons why IPv6 is the obvious choice for smart spaces. With IPv4 it 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, and 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, but this would not be recommended for a smart space.

Throughout the history of the Internet there has been a debate about the potential benefits of end-to-end communication [Clark et. al. 2002], [Reed et. al. 1998], [Saltzer et. al. 1984]. 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 (Network Address Translation). 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. The TSSG is a partner in a new EU Sixth Framework IST project: SEINIT (Security Expert Initiative). This project started on 1st December 2003 and has an ambitious schedule of research into security issues in IPv6.

Theoretically IPv6 promises easier provisioning of QoS (Quality of Service). This may or may not happen. 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 telecommunications networks. Smart space frameworks could allow for the potential for IP-based QoS 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] 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 through there is a critical mass of deployment of IPv6, particularly in Japan and Korea, but also in the cluster of EU funded projects in the Fifth Framework [6LINK 2002] [6LINK 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. 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).

4. Discussion & Conclusions

This paper has argued that smart space 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 from StockholmOpen.net.

The paper has also argued that the case is clear for the use of IPv6 as the network protocol of choice for smart spaces. Whilst there are still some potential issues, especially relating to security, with the use of IPv6, the advantages for mobile networks with large number of potential devices are so clear that the decision is obvious.

The paper has not addressed the higher level software infrastructure. 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, though one was retrofitted later REST (Representational State Transfer) [Fielding et. al. 2002]. The principle is that loosely coupled elements on a network could scale to provide useful services. Similarly it is possible to build smart spaces using lightweight architectures like the web itself (though even an embedded web server may be too heavyweight for some sensor devices). Perhaps it will be possible to use the next generation of Internet services, based around web services concepts (SOAP, WSDL, UDDI) to construct simple lightweight RESTful smart spaces over IPv6.

References

- [6LINK 2002] IPv6 Research and Development in Europe (November 2002) (on-line version available at <http://www.ist-ipv6.org/> [Last visited 2003-12-07])
- [6LINK 2003] IPv6 Cluster: Moving to IPv6 in Europe (July 2003) ISBN 3-00-011727-X (on-line version available at <http://www.ist-ipv6.org/> [Last visited 2003-12-07])
- [Battiti et. al. 2003] Roberto Battiti, Renato Lo Cigno, Fredrik Orava, Björn Pehrson "Global Growth of Open Access Networks:

from WarChalking and Connection Sharing to Sustainable Business" WMASH'03, September 19, 2003, San Diego, California, USA.

- [Clark et. al. 2002] David D. Clark, John Wroclawski, Karen R. Sollins, Robert Braden "Tussle in cyberspace: Defining tomorrow's Internet" SIGCOMM'02, August 19-23, 2002, Pittsburgh, Pennsylvania, USA.
- [Fielding et. al. 2002] Roy Fielding and Richard Taylor "Principled design of the modern Web architecture" ACM Transactions on Internet Technologies 2, 2 (2002) pages 115–150.
- [Intermon] <http://www.ist-intermon.org> EU FP5 IST Project Website [Last visited 2003-12-07]
- [Ó Foghlú 2002a] Mícheál Ó Foghlú "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
- [Ó Foghlú 2002b] Mícheál Ó Foghlú, Shane Dempsey, Eamonn de Leastar "Peer-to-peer Innovations for eBusiness and eWork: A Vision of Emerging Software Service Technologies" in Challenges and Achievements in E-business and E-work Edited by Brian Stanford Smith, Enrica Chiozza & Mireille Edin (Volume 2) IOS Press, Amsterdam, 2002 ISBN 1-58603-284-4
- [Reed et. al. 1998] David P. Reed, Jerome H. Saltzer, and David D. Clark. "Comment on Active Networking and End-to-End Arguments." *IEEE Network* 12, 3 (May/June 1998) pages 69–71.
- [Saltzer et. al. 1984] Jerome H. Saltzer, David P. Reed, and David D. Clark. *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.
- [Weiser 1993] Mark Weiser, "Some Computer Science Issues in Ubiquitous Computing," *Commun. ACM* 36, 7 (July 1993) pages 75–84.