Project Number:	**IST-2001-32603**
Project Title:	**6NET**
CEC Deliverable Number:	**32603/TERENA/DS/7.6/A1**
Contractual Date of Delivery to the CEC:	30 June 2005
Actual Date of Delivery to the CEC:	8 June 2005
Title of Deliverable:	Report on 3rd 6NET Open Workshop
Work package contributing to Deliverable:	WP7
Type of Deliverable*:	R
Deliverable Security Class**:	PU
Editors:	Kevin Meynell
Contributors:	Renzo Davoli, Patrick Grossetete, Latif Ladid, Stefano Lucetti, Kevin Meynell, János Mohácsi, Mario Morelli, Gabriella Paolini, Jean-Pierre Rombeaut, Marco Sommani & Gunter Van de Velde

* Type: P - Prototype, R - Report, D - Demonstrator, O - Other

** Security Class: PU- Public, PP – Restricted to other programme participants (including the Commission), RE – Restricted to a group defined by the consortium (including the Commission), CO – Confidential, only for members of the consortium (including the Commission)

Abstract:

This document provides a report of the 3rd 6NET Open Workshop held on 11-12 May 2005 in Pisa, Italy.

Table of Contents

Introduction

The 3rd 6NET Open Workshop was held on the 11th and 12th of June 2005 in conjunction with the GARR Conference 2005 in Pisa, Italy. The objective was to publicise the experiences of the 6NET project, and those of other related IPv6 developments. It also provided an opportunity for feedback from the Italian and wider European research networking communities.

The workshop was held over two days, in the form of one plenary session during the GARR Conference, followed by two dedicated workshop sessions. The focus was on making the case for using IPv6, drawing on the real-world experiences of the 6NET and Euro6IX projects amongst others. It also considered the benefits offered by IPv6 in the areas of multicasting and mobility, and presented an analysis of security issues.

In addition, the workshop covered the deployment of IPv6 in environments such as homes and schools which can benefit greatly from the improved multicasting and mobility functionality offered by the technology, but which also present a challenge with respect to security issues. The general approach should be that the average user need not involve themselves with the underlying technology, but they should be aware of the benefits that it brings.

All sessions were in English, although the final session focused on issues of specific relevance to the Italian research networking community. Presentations and approximate attendance figures were as follows:

11 June, 14.00-15.30 Chair: Marco Sommani Attendance: 150

- Enabling the two-way Internet and end-to-end innovation – *Latif Ladid, IPv6 Forum*
- Scaling the Internet for our Next Generations – *Patrick Grossetete, Cisco Systems*

11 June, 16.00-18.00 Chair: Gabriella Paolini Attendance: 53

- Mobile Wired Convergence and IPv6 – *Jean-Pierre Rombeaut, Alcatel*
- Euro6IX project and Italian IPv6 Task Force – *Mario Morelli, Telecom Italia Lab*
- Network Architecture Protection – *Gunter Van de Velde, Cisco Systems*
- IPv6 Security: Threats and solutions - *János Mohácsi, NIIF*

12 June, 09.00-10.30 Chair: Gabriella Paolini Attendance: 47

- Italian experience in 6NET – *Gabriella Paolini, Consortium GARR*
- IPv6 e mobilità - *Stefano Lucetti, University of Pisa*
- Multicast IPv6 - *Marco Sommani, CNR-IIT*
- An IPv6 Laboratory - *Renzo Davoli, University of Bologna*

The full proceedings of the workshop can be found on the 6NET website at:

http://www.6net.org/events/workshop-2005/

Enabling the two-way Internet and end-to-end innovation
Latif Ladid, IPv6 Forum

There are a large number of efforts around the world to promote IPv6. The IPv6 Forum which coordinates this is comprised of 180 organisations at the forefront of IPv6 development, including vendors, ISPs, research and academic institutes and software developers. It also has 10 chapters and 20 task forces targeting various countries and regions.

This effort has not only increased awareness of IPv6 in the technical community, but has increasingly worked to gain political support for IPv6 initiatives in many countries. This has been undertaken through advocacy, high-profile showcases and media coverage, not to mention a number of government-sponsored industry trials. Allied to this is the introduction of the 'IPv6 Ready' programme that marks products as being suitable for use with IPv6, with the aim of improving consumer awareness. A number of Global IPv6 Summits (twelve in 2004) are also organised in different regions, and with a total of 24,000 participants, they constitute a large informal IPv6 university.

There are currently a number of ISPs and NRENs offering IPv6, and it is already supported by major operating systems such as Windows, Mac OS X and Linux. It is also being adopted by the consumer electronic industry, and will be included in forthcoming PDAs, gaming consoles, webcams, and 3G mobile phones. It is also being adopted by the automobile industry for use in vehicles, and in the defence industry where the US Department of Defense has mandated its use for future procurements.

However, there is still some work to do. In particular, few of the 200 so-called point-to-point applications such as Skype support IPv6, whilst people need to be assured that the end-to-end addressability of IPv6 offers the same levels of security that IPv4 is perceived to. There is also the issue to be addressed that whilst IPv6 has been designed to be largely transparent to end-users, that invisibility makes it more difficult to persuade them to adopt the technology.

This presentation can be found on the web at:

http://www.6net.org/events/workshop-2005/ladid.pdf

Scaling the Internet for our Next Generations
Patrick Grossetete, Cisco Systems

There is an increasing trend of convergence between data, voice and video services, and the services available across different underlying infrastructures. This is combined with advances in bandwidth, mobility, security and multimedia that make networks even better for the consumer. IPv6 is needed to help deliver this, but to network managers it must be proven to be stable and offer easier deployment than existing IPv4 technology. To the average end-user, they only care about what services and applications it can deliver and do not wish to be bothered with the underlying technology.

Revenues from traditional analogue voice and bandwidth services are steadily decreasing, so there is a need for a new model. In particular, there is a need for a different address allocation policy and charging model, because allocating more addresses does not necessarily mean that networks will

become over-subscribed. ISPs can gain additional revenues from providing services to end-points, rather than merely providing bandwidth. In particular, there is a great deal of scope for increased use of network-enabled devices in homes (e.g. VoIP, gaming, video monitoring, video-on-demand), and also for mobile personal devices that can roam between networks.

Of course, this will also lead to changes in network traffic. Traffic is likely to become more symmetrical as users start serving content as well, sessions will become longer as always-on devices are left unattended, multiple video sessions can sustain high bandwidth, whilst more traffic will travel globally putting load on peering points and expensive long haul links.

IPv6 enables the Internet to be expanded, not only with respect to the number of addresses, but also though the deployment of scalable technologies that are associated with it. The core IPv6 specifications are already stable and have been implemented in production networks such as 6NET and various NRENs, and they offer possibilities for new uses in the coming years. For example, IPv6 has already been trialled in the automotive industry for in-vehicle navigation and traffic information systems, whilst integration of different technologies such as DECT and 3GPP becomes possible. IPv6 also offers opportunities for schools which are frequently forced to use a single IPv4 address with NAT, limiting the possibilities for remote learning, inter-school collaboration, secure information exchange and tele-surveillance.

All Cisco IOS based networks have been IPv6 capable since 2001, and the solutions offered now include routers, Layer 3 switches, firewalls and network management facilities. Cisco is heavily involved in the IETF standardisation process, and co-chairs several of IPv6-related working groups. It was also a founding member of the IPv6 Forum, and has established partnerships with a number of large-scale IPv6 deployments such as 6NET and Moonv6.

This presentation can be found on the web at:

http://www.6net.org/events/workshop-2005/grossetete.pdf

Mobile Wired Convergence and IPv6
Jean-Pierre Rombeaut, Alcatel

Fixed Mobile Convergence (FMC) is becoming increasingly necessary for the consumer who may have a wireless DECT telephone at home and a GSM mobile telephone for elsewhere. Trials in Japan have demonstrated that it is possible to have the same telephone for multiple network access, and the increasingly ubiquitous nature of IP makes this the natural choice for converging the various technologies.

The 2006-2007 timeframe will be critical to FMC technology and services development, with boundaries between fixed and mobile technologies gradually disappearing by 2010. Service providers are taking an optimistic view of FMC, believing that it will bring fundamental changes to the structure of the telecom market.

There are several key technologies necessary for FMC. These include mobility which allows users to roam between different access points, but which has potential difficulties with incompatible authentication and registration mechanisms. They also include session control which manages connections originating and terminating in both the fixed and mobile environments. QoS is required

to ensure certain guarantees of session quality, whilst security ensures that only authorised users are able to use network resources. Finally, there needs to be a move away from standardised services towards standardising service capabilities which provide the flexibility to deploy and offer new services.

The different access technologies need to use a common protocol in order to be unified, and IP is the best candidate for this. Whilst IPv4 could in principle be utilised, IPv6 offers more flexibility as well as better support for Mobile IP. In particular, IPv6 uses IPSec, it offers autoconfiguration so handovers can (in theory) be faster, and MIPv6 is native to the IPv6 specification. MIPv6 is a way for an end-user to roam between different networks whilst maintaining a permanent address at which they can always be found. This can be used in conjunction with different access technologies (e.g WiMax and 3GPP) to provide seamless connectivity with the Internet.

A number of telecoms operators including BT, France Telecom, Swisscom and Korea Telecom will shortly be offering devices (e.g. phones, PC cards, software clients) that provide wired/wireless convergence.

This presentation can be found on the web at:

http://www.6net.org/events/workshop-2005/rombeaut.pdf

Network Architecture Protection
Gunter Van de Velde, Cisco Systems

One of the difficulties in persuading users to adopt IPv6 is that the widespread use of Network Address Translation (NAT) is perceived to provide a degree of privacy and protection to IPv4 networks. NAT hides end systems from the wider network and provides simple security due to a lack of translation state. It also allows for addressing autonomy and easy renumbering, as translation is undertaken at the interface between private and public networks.

The primary purpose of NAT is to increase the available address pool, but this is not necessary for IPv6 which has more than enough addresses for the foreseeable future. NAT also has a number of well documented disadvantages, but its perceived market benefits means that IPv6 must be seen to offer similar levels of protection. This is why it was felt necessary to produce an Internet Draft to document the techniques that may be combined on an IPv6 site to protect the integrity of its network architecture.

There are situations where it is desirable to prevent device profiling, such as when web sites are contacted, so IPv6 privacy addresses were defined to provide that capability. IPv6 addresses typically include an interface identifier based on a MAC address which in practice facilitates tracking of a device, but RFC 3041 provides a privacy address extension that randomises the interface identifier. However, this only precludes tracking on the lower 64 bits of an IPv6 address, so something like a Mobile IPv6 gateway is required to obscure the destination of globally routable prefixes.

Another alternative is Unique Local Addresses (ULAs) which are prefixes set aside for local communications. They allow networks to be privately connected without address conflicts or renumbering, and provide a known prefix that can be filtered at network boundaries. This allows

routers to establish routes between local networks and the global Internet, but without revealing internal topologies.

The vulnerability of an IPv6 host is similar to that for an IPv4 host directly connected to the Internet, so firewall systems are recommended. However, IPv6 by default offers short lifetimes on privacy extension suffixes, and IPSec functions to prevent hijacking and content privacy. Whilst IPSec is available for IPv4, this is generally not compatible with NAT. In addition, the sheer size of a typical IPv6 subnet (::/64) makes a network ping and/or port scan somewhat impractical.

Multihoming and renumbering still remains problematic under IPv6, but the protocol was designed to allow sites and hosts to run with several simultaneous prefixes and hence ISPs. A renumbering effort may be required if a network changes provider, but with appropriate management of the prefix update mechanisms, a smooth transition can be effected.

There are still some gaps in the standardisation process, namely the finalisation of the ULA format, the need for a topology masking component for RFC 3041 addresses, and documented renumbering/multihoming procedures (which to some extent require ULAs). There should also be some work with Mobile IP to allow a home agent to use internal tunnelling to reach the mobile node, thus never revealing its real location.

This presentation can be found on the web at:

http://www.6net.org/events/workshop-2005/vandevelde.pdf

Euro6IX Project and Italian IPv6 Task Force
Mario Morelli, Telecom Italia Lab

Euro6IX is an IST funded project that started in 2002 and involves the research departments of seven telecommunications operators, as well as number of universities. The goal was to design and build an IPv6-only infrastructure to test IPv6 services and applications.

The network featured seven nodes connected by 34 Mbps links, each of which was Internet Exchange. It was intended to be a multi-vendor trial, and used equipment from Cisco, Juniper, Hitachi and 6WIND amongst others. There were also interconnections with the 6NET and Moonv6 networks, and connectivity with the virtual 6Bone.

The traditional model of an Internet Exchange is the provision of a Layer 2 high-performance infrastructure where ISPs come to exchange their traffic. No end-user services, IP addresses or even Layer 3 services are usually provided. The problem is that IPv6 prefixes are provider dependent and if an end-user changes provider, then their addressing space also has to be changed. However, RFC 2374 proposes an addressing scheme where certain network points (e.g. IXs) allocate addresses to their customers so that renumbering is not required.

Euro6IX has therefore investigated different IX models based on the concept that it was an aggregation point for both ISPs and large customers. If services (e.g. AAA, DNS multicast and route servers), and applications (e.g management and security) could be placed inside the IXs, everyone coming to the IX would be able to benefit from aggregation and improved performance. If

IXs actually provided addresses, that would make some operations like renumbering and multihoming easier to implement.

The project also investigated a Policy Based Management Network (PBMN) tool which could centrally manage and monitor policy inside an IX, and a route server that could handle RPSLng database implementations. In addition other applications were developed including the Magalia monitoring tool, the ISABEL conferencing suite, a SIP audio client and an IPv6 mobility demonstrator.

The presentation concluded with an overview of the Italian IPv6 Task Force that had been formed in October 2003. This involved thirty partners including ISPs, telcos, IXs, vendors, universities and research centres, and had the goal of spreading awareness of IPv6 in Italy. It usually met twice per year, and had five different working groups in the areas of public administration, private industry, mobility and wireless, dissemination, and business models.

This presentation can be found on the web at:

http://www.6net.org/events/workshop-2005/morelli.pdf

IP Security: Threats and Solutions
János Mohácsi, NIIF/HUNGARNET

The introduction of IPv6 brings with it a number of new security issues that must be taken into consideration. As the default subnets in IPv6 have 2^64 addresses, an exhaustive scan on every address is no longer practical. However, some hosts will still need to be in the DNS, administrators may adopt easy-to-remember addresses (e.g. ::1,::2,::53), EUI-64 address have a fixed component, and it may be possible to guess addresses from the MAC addresses used by vendors. The new multicast addresses can also enable attackers to identify key resources on a network, unless they are filtered at a border point in order to make them unreachable from outside.

Many current IPv4 networks make use of Network Address Translation (NAT) which hides end addresses and acts as a rudimentary firewall. IPv6 does not require NAT as it allows for end-to-end addressing, but this means that firewalls become even more important. End-to-end addressability does not necessarily mean that all addresses have to be reachable from the outside world, although security measures become more complex when both IPv4 and IPv6 are being run together.

The threats to an IPv6 network include header manipulation and packet fragmentation, Layer 3/4 spoofing (which is easier because IPv6 addresses are aggregated), ARP cache poisoning, DHCP snooping, and BGP and IS-IS routing attacks. Of course, IPv6 networks are also vulnerable to traditional viruses and worms, packet sniffing, and flooding attacks as well. Nevertheless, there are a number of IPv6-enabled firewalls available that can protect against such attacks, although few currently support FTP and H.323 very well.

When running a mixed IPv4 and IPv6 environment it is very important to ensure the same security mechanisms are implemented for both protocols. In addition, a network administrator must be very careful about tunnels as these can easily bypass any firewall policies. Completely automated tunnels should therefore never be used, translation mechanisms should be avoided, and only authorised systems should be allowed as tunnel end-points. Where possible, IPSec should be utilised (even

though it is supposed to mandatory in IPv6) as this provides good authentication and confidentiality.

This presentation can be found on the web at:

http://www.6net.org/events/workshop-2005/mohacsi.pdf

Italian Experience in 6NET
Gabriella Paolini, Consortium GARR

The Italian academic and research community has been at the forefront of IPv6 implementation, initially through the 6Bone, and though involvement with the IETF and RIPE. GARR had also been a key partner in the 6NET project which had operated the largest native IPv6 network in Europe in the world. Ten Italian universities and research institutes were directly involved in this project, and were connected by an IPv6 network with seventeen nodes. This was provided via IPv6 over ATM, or IPv6 tunnels over existing IPv4 links, running OSPFv3.

Several IPv6 training workshops had already been held in Turin, Rome, Florence, Bari, Milan, Naples and Cagliari, with three more due to be held in Genova, Messina and Trieste. Material from these workshops was available on the 6NET Italia website at http://www.6net.garr.it/

This presentation can be found on the web at:

http://www.6net.org/events/workshop-2005/paolini.pdf

IPv6 e mobilità
Stefano Lucetti, University of Pisa

The Mobile IPv6 protocol (MIPv6) is an IETF standard developed to provide transparent mobility support to hosts within IPv6 networks. MIPv6 allows a Mobile Node (MN) to move across different networks, keeping its own identity (in the form of its static, persistant Home Address - HoA), whilst being reachable by means of a temporary address (indicated as Care-of Address - CoA) belonging to the visited network.

The protocol intelligence resides in both the Mobile Node itself and the Home Agent (HA), a new entity introduced by MIPv6 which belongs to the home network of the MN rather than the network core. The HA takes care, among the other things, of keeping an up-to-date association between the HoA and the current CoA of the MN (if away from the home network), as well as of forwarding packets addressed to the HoA of the MN to the current CoA by means of IPv6 encapsulation. Furthermore, MIPv6 incorporates a route optimization strategy to allow direct communication between the MN and its Correspondent Nodes (CNs), as well as proper security mechanisms to protect signaling among MIPv6 peers.

From a practical point of view, MIPv6 enables a MN to be univocally identified by a well known IPv6 address (it may be the one recorded in DNS), independently of its current point of attachment to the network, and hence of its geographical position. It is noteworthy that such property is not related to a particular interface, but applies also to the case where the MN changes the interface through which it accesses the network. This inter-technology handover is often indicated as Vertical

Handover, and is particularly interesting if the current trend of network access towards overlay approach is considered. Here, the term overlay means that several technologies concurrently offer network access at the same location.

MIPv6 allows the MN to switch among the different access networks (characterized by different advertised IPv6 prefixes) transparently to applications, keeping alive its ongoing transport level connections. A predetermined or dynamically evaluated priority value is associated to each interface, so that higher-data rate, low-cost wired and WLAN connections are preferred to WWAN ones. Typically, Ethernet interface has the higher priority, followed by WLAN and WWAN ones, respectively.

The demonstrator developed at the University of Pisa aims at show both these mobility philosophies are supported. The test-bed is composed of four IPv6 subnets, interconnected through an IPv6 network emulated by a single IPv6 router. One of the subnets is the home network of the MN, where the HA resides; whilst a generic CN is hosted in a second subnet. The remaining two subnets act as visited networks for the MN; one is a wired Ethernet, representing a remote network, whereas the other is a Wireless LAN based on 802.11g, which overlays to the others. A DNS server located at the HA provides name resolution for the demo6net.org local domain used for the demonstration.

Two different, although strictly related, experiments are shown. The first is oriented to the demonstration of MIPv6 capability to maintain the identity and reachability of a MN located in a remote visited network. The second aims to highlight the transparent rerouting of active connections through different interfaces after a vertical handover occurs.

In the first experiment, the MN is initially connected to its home network, but then moves to the remote wired one (no active WLAN access is active during the handover) where it acquires a CoA and registers it with the HA by sending a Binding Update message. The CN queries the MN at its HoA, and the HA intercepts the packets and redirects them to the CoA of the MN using IPv6 encapsulation. As soon as the MN receives the first tunnelled packet, it starts the Return Routability procedure in order to authenticate the successive Binding Update transmitted to the CN to perform the Route Optimization. From that moment on, all the communications are directly between the MN's CoA and CN, taking advantage of the extension headers introduced by MIPv6.

The second experiment focuses on vertical handover and its performance. The MN, at one of the wired subnets, has active TCP sessions and/or UDP unicast flows with the CN. Upon a predetermined expiration time, the MN gets disconnected from its home network and the MN registers the CoA associated with the WLAN subnet to its HA, All ongoing connections are then transparently rerouted through the wireless interface. This demonstrated with a variety of different applications including file transfers using FTP and streaming of multimedia flows.

This presentation can be found on the web at:

http://www.6net.org/events/workshop-2005/lucetti.pdf

Multicast IPv6
Marco Sommani, CNR-IIT

IPv6 offers improved multicast facilities over IPv4, with its scoped addressing and embedded Rendezvous Points (RPs). Unlike IPv4 multicast that uses special Class D addresses that must be handled separately, IPv6 can include scoping and group identifiers as an integral part of the address. These are identified through pre-defined flag bits.

There are two main types of multicast routing. Any Source Multicast (ASM) uses IGMPv2 and MLDv1, and allows multiple hosts to send traffic within a multicast group. This requires a complex group and tree setup process, and can result in excess network traffic and collisions. By contrast, Single Source Multicast (SSM) learns about sender IP addresses from applications rather than RP routers. This makes it much simpler and results in more reliable multicast sessions, but it relies on IGMPv3 and MLDv2 which is not currently supported by many routers or applications.

With SSM, different groups can use the same multicast addresses with no risk of collisions, but with ASM it is very important to ensure the uniqueness of the multicast addresses. IPv6 provides two ways to solve this problem: unicast prefix-based addresses and embedded RP addresses.

Another problem with ASM is how to allow the coexistence of multiple RPs in the Internet. IPv4 had found a solution with the MSDP protocol, but this has scalability problems and can quite easily become the victim of DoS attacks. In IPv6 though, Embedded RP introduces an algorithm for deriving the RP address from the multicast address, and simplifies management of the network. This is not possible with IPv4 as the 32-bit addresses that it uses are simply not large enough to incorporate the RP addresses.

The 6NET project did a lot of work on IPv6 multicast deployment, particular with respect to Embedded RP. This is documented in several deliverables (D3.1.2v2, D3.4.2 and D3.4.3) and a number of Internet Drafts that are available at http://www.6net.org/.

This presentation can be found on the web at:

http://www.6net.org/events/workshop-2005/sommani.pdf

An IPv6 Laboratory
Renzo Davoli, University of Bologna

The University of Bologna is running an IPv6 laboratory which is connected to both the existing IPv4 network and the Italian IPv6 testbed. It is useful to have students utilise IPv6 because they generate a lot of traffic and they are demanding users.

A number of interesting tools are being used to facilitate the laboratory. Prometeo provides all the shared services needed by proxies (e.g. DNS caching), which can be dynamically created with loadable modules. Another useful tool is vish which is a programmable firewall used for authenticating users on both wired and wireless networks. This is used to provide public access services to student laptops. The dbind program provides dynamic DNS services for these laptops, as well as automatic name generation for large clusters. Finally, VDE (Virtual Distributed Ethernet)

offers Layer 2 overlay functionality that can provide connectivity to different kinds of software components including emulators, virtual machines, operating systems and other connectivity tools.

VDE is a distributed system that does not need specific administrative privileges to run. It consists of two components: the vde_switch which is the virtual counterpart of a physical Ethernet switch, and the vde_cable which is the virtual counterpart of a crossover cable. This allows different virtual topologies to be created and utilised, as well as providing enhanced security and privacy.

The next stage is to implement ViewOS which allows the semantics of system calls to be independently changed by users. Each process can then see different resources (e.g. file systems, devices or network resources). The current implementation called UM-ViewOS runs on unmodified Linux kernels using the ptrace interface for debuggers, and does not require root access. There is also a kernel patch that increases performance when installed, although this does not add any extra features.

This presentation can be found on the web at:

http://www.6net.org/events/workshop-2005/davoli.pdf