

Project Number:	IST-2001-32603
Project Title:	6NET
CEC Deliverable Number:	32603/Partner/DS/No./A1
Contractual Date of Delivery to the CEC:	June 2005
Actual Date of Delivery to the CEC:	June 2005
Title of Deliverable:	IPv6 Basic Network Services Whitepaper
Work package contributing to Deliverable:	WP3
Type of Deliverable*:	R
Deliverable Security Class**:	PU
Editors:	João Nuno Ferreira
Contributors:	Tina Strauf, Christian Schild, Gunter Van De Velde, Jan Novak, Jerome Durand, Stig Venaas, Christian Schild, Tina Strauf, Carlos Friaças, János Mohácsi, Sabine Kuehn, Ana Romero, Wilfried Woeber, Bruno Ciscato, Olivier Courtay, João Pagaimo, Alexander Gall, Bruno Ciscato, Bartosz Belter, Michał Balcerkiewicz, Ralph Droms, Francis Dupont, Bartosz Gajda, József Kadlecik, Marcin Kamiński, Blazej Pietrzak, André Stolze, Simon Leinen, Rob Evans, Gabriella Paolini, Georgios Koutepas, Athanassios Liakopoulos, E. Rosti, E. Vyncke, D. Zacharopoulos, Tim Chown, Mark Thompson, Alan Ford, Thorsten Kuefer, Frédéric Beck, Olivier Festor, Bartek Gajda, Kurt Bauer, Wim Biemolt, Piers O'Hanlon, Jean-Jacques Pansiot, Antonio Pinizzotto, Lorenzo Rossi, Tomasz Szewczyk, Ralph Droms, Bernard Tuy, Antonio Pinizzotto, Mickael Hoerd.
Reviewers:	Mónica Domingues, Mickael Hoerd

* 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:

A short and concise description of the IPv6 development status is given, focusing on the Basic Network Services that were included in Workpackage 3 of the 6NET project.

Keywords:

IPv6, Routing, Multicast, DHCP, RPSLng, Renumbering, Security.

Executive Summary

As the 6NET project approaches its end, a great deal of knowledge was accumulated in the project Deliverables. Most of this Deliverables are extensive and much focused on one or two particular IPv6 issues. It was recognized that for those interested in having a quicker reading about the present status of all those issues, it would be very useful to have a higher-level and shorter, set of documents. This Whitepaper attempts to fulfil that goal in respect to the Basic Network Services targeted in Workpackage 3 of the 6NET Project.

It was also considered important to produce a broader view of all the important achievements attained in IPv6, that sometimes are not adequately emphasised as all efforts continuously move on in the quest for solving the remaining open issues.

Table of Contents

1.	INTRODUCTION.....	5
2.	ROUTING.....	6
3.	DNS AND DHCPV6.....	6
4.	REGISTRY (RPSLNG).....	8
5.	MULTICAST	8
6.	SECURITY	9
7.	RENUMBERING.....	9
8.	CONCLUSION.....	11

1. Introduction

After more than three years of work, the 6NET project achieved several important goals. In the IPv6 Basic Network Services (BNS) area the progresses made during this period was remarkable. BNS are the building blocks of any network infrastructure. They remain mostly invisible for the end user, but it's where most of the times issues like complexity, scalability and interoperability come to play a central role.

Workpackage 3 (WP3) was divided in five sub-areas, to which a sixth was added later:

1. Routing
2. DNS and DHCP
3. Registry (RPSLng)
4. Multicast
5. Security
6. Renumbering

In this Whitepaper we followed the same division to present the main conclusion in each one of these areas. References to the corresponding deliverables were included at the end of each section.

2. Routing

In the first six months of the project an enormous effort has been made to build up the 6NET network. In 35 locations Cisco routers have been installed, the national, backbone and local loops have been ordered and tested; the routing and addressing plan has been implemented and deployed. In the second half year, with beginning of August, almost all NREN and Universities were connected to 6NET, based on native IPv6 connections.

In addition to the original 31 6NET partners, four new partners (PSNC, CESNET, Hungarnet and ETRI) joined the project on 1st September 2002, as well as FCCN later in 2003. To incorporate the new partners, the network has been extended with IPv6 tunnels to Poland, the Czech Republic, Hungary and Portugal. Moreover, the geographical coverage of the network and with it the scope of the activities has been broadened by connecting 6NET to other research networks in Europe, North America and Asia. 6NET provides connectivity for example to EURO6IX, 6BONE, Abilene and NTT. Although 6NET is a test-bed network, one of the major requirements was to keep the network stable. Special maintenance and test windows within the network core have been defined to realise that.

Apart from the operation of such a large continent-wide IPv6 network, a separate task was developed to gather implementation status knowledge on:

- Router Advertisements
- Neighbour Discovery
- RIPng
- IS-IS
- OSPF
- BGP

It was found that all these protocols are available today in several manufacturers in production stable version provided with maintenance.

The knowledge and confidence gained in 6NET enabled the anticipation of the launch of GEANT IPv6 Unicast Service, in October 2003. Since then this service has been operated in a dual-stack mode to all NRENs connected to it. In April 2004, Class-of-Service (CoS) was successfully implemented in the 6NET backbone, further enriching the complexity of the IPv6 routing function.

We are confident in reaching the conclusion that Routing can be regarded as one of the most well developed, tested, interoperable and available IPv6 areas.

For more information on the design and operation of the 6NET backbone, see Deliverables D3.1.1 and D1.5.x.

For more information on configuring Routing in several equipments, see Deliverable D3.1.2.

3. DNS and DHCPv6

At the time of 6NET planning DNS and DHCP were grouped together. Later it became evident that most of the work was to be done in the DHCP part rather than the DNS one.

3.1. DNS

The DNS service is a corner stone of any IPv4 or IPv6 network. As IPv6 uses the same hierarchy as IPv4, the dual-stack behaviour of this service is central to have a smooth operation and integration of both protocols.

After the setup and operation of the first set of services early on the first year of the project, namely DNS forward and DNS reverse, it was envisaged that DNSSEC would be added in the 6NET backbone. But this suffered a considerable delay regarding its initial plan. Two set of problems caused this. Firstly, although it was initially included in the Description of Work, as the project progressed it became clear that the usefulness of dealing with DNSSEC issues in an IPv6 project like 6NET was questionable.

Secondly and more important, the DNSSEC standard that was stable at the beginning of the project proved in several trials to have severe practical implementation problems and soon the work on an alternative and better standard started. The schedule of this second standard, which is incompatible with the previous one, slipped forward in time, until by the end of 2003 when the Project Review was done it was still not available.

The decision was then to either wait until new and updated DNSSEC software became available or write down these difficulties together with an implementation plan for a future implementation of DNSSEC in 6NET. The second option was chosen and when the new DNSSEC standard implementation was available, in the second quarter of 2004, it was installed and successfully tested in the 6NET backbone. That proved the expected isolation between DNSSEC and IPv6.

Most important was the decision to allow IPv6 glue records for ccTLD in the root-servers. Soon after this became possible, in the 2nd quarter of 2004, dozens of ccTLD asked for the insertion of these records. This enables the creation and operation of fully compliant IPv6 DNS domains.

For more information on configuring DNS see Deliverables D3.1.2 and D3.2.1.

3.2. DHCPv6

After many years of work, DHCPv6 was published as RFC 3315. Despite the existence of stateless autoconfiguration for IPv6 (RFC 2462), there is still a need for DHCP. On one hand, it complements stateless autoconfiguration where it can supply hosts with DNS, NTP and other configuration data. On the other hand, a network administrator might want to gain more control over the IP addresses used than what is possible with stateless address configuration. A stateful DHCPv6 implementation as of RFC 3315 offers both. In addition, the IETF DHC working group published a more lightweight respective "stateless" DHCPv6 version (RFC 3736), which serves only as a source for configuration options that are not already delivered to the host with stateless autoconfiguration.

The stateless mode of DHCPv6 and especially the prefix delegation and nameservice option are more widely implemented as they seem to be more important for today's networks. To the knowledge of the authors prefix delegation and nameservice option should work fine for most implementations. The lack of DHCPv6 relays is not so important in the stateless mode, as one could set up a DHCPv6 server on every link's router. So at least the stateless feature sets seem to be deployable today.

Speaking of address delegation with DHCPv6, the implementations mostly seem to be immature. In addition, the implementations do not interoperate properly. Obviously there has to take place more interaction between the different developers. So from a today's point of view, stateful DHCPv6 service is not deployable or at least only with some limitations. But there were giant steps since the beginning of 6NET. So it seems that there will be a complete implementation in the near future for the major operating systems.

For more information on configuring DHCPv6 see Deliverable D3.2.3.

4. Registry (RPSLng)

The Internet is the interconnection of many independently managed networks. Each network is a domain with its own policies about internal and external relationships. The protocol used to exchange routing information between domains, also known as «Autonomous Systems» (AS), is the Border Gateway Protocol version 4. The Multiprotocol Extensions to BGP (MP-BGP) are widely implemented and deployed today, mainly for IPv4 multicast, IPv6 unicast, and BGP/MPLS VPNs. Deployments for IPv6 multicast are also starting to show up, mainly in the Research & Education Networks' environment.

The exchange of BGP routes between ASes is constrained by policies. Operators have the possibility to document their routing policies in their regional routing registry. In Europe the Regional Internet Registry is RIPE, but alternative routing data repositories also exist (RADB is an example). The language, that standardizes entries format in routing registries, has been defined in RFC 2622 - Routing Policy Specification Language and in RFC 4012 - RPSLng.

The information in the Routing Registries can be used, above all, for automatic generation of BGP router configuration. Policy information can be documented in the registry at various levels of detail. The operators decide its policies and their level of publicity. Most upstream providers ask their customers to maintain routing registry data continuously up to date and consistent. Conditions can be relaxed in some Internet Exchange Points, especially for small organizations.

After a long and slow development and testing phase, RPSLng has become a standard and it is active in the RIPE database. Although few operators have practical experience with the proposed RPSLng specification, it can be observed a growing number of RPSLng policies, as well as IPv6 addresses assignments by the LIRs to their end-users.

Like IPv4 Routing Policies, IPv6 routing policies need to be documented. Maintaining up-to-date routing policy information data should be a daily practice, but unfortunately often network administrators do not give it high priority, or decide that is better for them not to completely show their policy. The continuous maintenance of RPSLng tools projects should be a substantial contribution to the future implementation that we hope for the new version of routing policy language.

For more information on configuring RPSLng see Deliverable D3.3.1

5. Multicast

Using unicast could result in sending the same message as many times to as many recipients there are. While using multicast, the message would be sent only once from the sender to the multicast group, the interested listeners. To achieve this, routers in a network, build a distribution tree. The advantages of multicasting are bandwidth demand reduction, server load decrease and less hardware requirements.

Although reliability and congestion control can be built on top of UDP (and hence above the transport layer), there is still no generic mechanism and most multicast applications do not provide any congestion control today (which can be a real problem if multicast is widely used).

Multicast was one of the 6NET areas with development work active until the end of the project. The project partners made a proposal for a new standard for IPv6 Inter-Domain ASM: The Embedded-RP Solution, that's already available in several router manufacturers. Some issues remain open with Embedded-RP, especially if the service is deployed in core networks, namely: Group address allocation, RP access control, RP robustness issues and RP securities issues.

With the help of the expertise acquired, GEANT2 launched its pilot Multicast Service in the 1st quarter of 2005.

A multicast addressing study was also produced as well as a Deliverable regarding IPv4/IPv6 Multicast interoperability.

For more information on Multicast, see Deliverables D3.4.x.

6. Security

IPv6 is regarded worldwide as a solution for a number of the networking problems, the most important of which being IP address depletion. It also provides additional features that can enhance the network interconnection characteristics and security. Many of these new features however, have been untested in a real networking environment and could present novel challenges and threats.

6NET started as a network for providing core IPv6 connectivity, between European academic institutions. Since its successful implementation the focus has shifted from the basic (core) networking services to user connectivity and services. It's these services that provide the real value of the network and usefulness to the end users. Although in 6NET the assets at risk are academic networks and the operation of experimental services, security problems can escalate to the universities' internal networks, possibly creating problems there, but also may disrupt experiments and other evaluation processes of 6NET itself. Furthermore, since 6NET is also an important feasibility demonstration project, security will guarantee the public's trust in the value and especially reliability of the new technology.

One of the characteristics of 6NET that adds importance to security issues is its high speed network connections. Any high bandwidth line offers a very efficient infrastructure for the malicious users that will manage to utilize it to attack any of its nodes or targets outside. Specifically in IPv6, 6to4 facilities provide the possibility of an attack "spill-over" to more critical IPv4 production and even commercial networks. Policies, methodologies, configurations and the general security lessons learn from 6NET will, hopefully, prove to be a useful and substantial contribution to the IPv6 community and future deployments.

The target audience for this knowledge is network managers and administrators that control medium to large size networks and want to transition to IPv6, understanding the threats and requirements of the new environment, without compromising their security level. 6NET Deliverables offers them some guidelines on the security issues they may encounter.

The vastness and fragmentation of the security issues prevents it from having a concise set of conclusions which could be included in this whitepaper. As such we strongly stress those interested in reading Deliverable D3.5.1v3.

7. Renumbering

An IPv6 site will need to get Provider Assigned (PA) address space from its upstream ISP. Because Provider Independent (PI) address space is not available for IPv6, a site wishing to change provider will need to renumber from its old network prefix to the new one.

Scenarios, issues and enablers for such renumbering were considered. There are a number of conclusions that can be drawn from the work done. These include:

- Tools and methods to ease renumbering of an IP network have not been advanced as much as they might have been, most likely because IPv4 offers Provider Independent (PI) address space for large enterprises. The PIER WG began studies in 1996 (or so) but closed shortly thereafter after publishing two RFCs.
- There are some features of IPv6 that lend themselves to easing renumbering (stateless autoconfiguration, DHCP-PD, Router Renumbering) but these still need further deployment trials. In addition, many of the common IP problems remain (e.g. hard-coded IP literals, in local and remote systems).
- Scenarios and triggers for network renumbering appear to be relatively well defined and understood.
- A phased process for IPv6 renumbering without a flag day has been proposed. Trials with ‘simple’ SOHO scenarios suggest the procedure is sound. Further work is required in larger enterprise scenarios, where DHCP-PD may assist, for example.
- Network backbone renumbering following was generally sound, but found some issues, e.g. in multi-hop BGP usage (and ordering of changes). The IPv6 renumbering features did not really help with the backbone renumbering process (it seemed as complex as IPv4 renumbering).
- Network management tools (still in their relative infancy for IPv6) are generally not well adapted to handle IPv6 renumbering events. The tools need restarting, statistics files need to be ‘post-processed’, or other undesirable properties would manifest themselves, e.g. multi-addressed devices might appear as multiple nodes on the network. There is an impact on both management and monitoring applications.
- Running an IPv6 link with two (non link-local) prefixes active is quite possible, but a number of issues have been identified in this text. Further experimentation is desirable.
- A set of recommendations has been identified for key target audiences (network designers and administrators, application developers, and operating system/stack developers). These need to be disseminated to those audiences for discussion and progression where possible. Further recommendations may also be drawn.

For more information on Renumbering, see D3.6.1.

8. Conclusion

In the course of the 6NET project, and in particular, in WP3 all the major BNS were studied, tested and deployed either in the separate 6NET backbone, or in the GEANT(2) backbone. The main conclusions are that BNS are readily available today from manufacturers, making it possible to build a production IPv6 network without any particular difficulty or barrier. Those areas, where further development was needed, made huge progress during the project life and are now being completed.

Today, IPv6 is a smooth running technology, with unparallel scalability, and with a growing number of providers and users.