

Project Number: IST-2000-26417
Project Title: GN1 (GÉANT)



Deliverable D9.3

IPv6 Testing

Deliverable Type: PU-Public
Contractual Date: 30 June 2001
Actual Date: 14 August 2001
Work Package: WI8.4
Nature of Deliverable: RE - Report

Authors:

Text compiled and edited by Tim Chown (University of Southampton and UKERNA, UK).
See Acknowledgements for test participant list.

Abstract:

In this report we describe the IPv6 tests that have been performed by participants in the IPv6 Working Group within the GÉANT Task Force for Next Generation Networks (TF-NGN). The group identified a number of IPv6 work items. These were studied from a deployment-led perspective, including a testbed network through which a core router connected up to fifteen participants during the course of the work. The results of the tests were generally positive, indicating that the basic building blocks are in place for moving towards a production IPv6 deployment. Future work within the TF-NGN IPv6 WG will continue to pilot IPv6 activity, and also to consider and report on IPv6 work items that were only given secondary priority during the current phase of reporting.

Keywords:

IPv6, IPv4-IPv6 Transition, DNS

CONTENTS

CONTENTS	2
ACKNOWLEDGEMENTS	3
EXECUTIVE SUMMARY	4
1 INTRODUCTION	5
2 PRIMARY WORK ITEMS	7
2.1 PLATFORMS, ROUTING AND INTEROPERABILITY	7
2.1.1 <i>Current GTPv6 core IPv6 Network</i>	7
2.1.2 <i>Evolving the GTPv6 core network</i>	9
2.1.3 <i>IPv6 Support in Routers</i>	11
2.1.4 <i>UNINETT (Norway)</i>	11
2.1.5 <i>UKERNA (UK)</i>	12
2.1.6 <i>RedIRIS (Spain)</i>	14
2.1.7 <i>JOIN / DFN (Germany)</i>	15
2.1.8 <i>CESNET (Czech Republic)</i>	18
2.1.9 <i>POZNAN (Poland)</i>	18
2.2 DNS	20
2.2.1 <i>DNS activity by partners</i>	20
2.2.2 <i>Test results</i>	21
2.3 REGISTRIES AND ADDRESSING	22
2.4 TRANSITION TOOLS	23
2.4.1 <i>Partner activities</i>	24
2.5 APPLICATIONS	25
2.5.1 <i>TIPSTER6 Applications Database</i>	25
2.5.2 <i>Application and porting work</i>	29
2.5.3 <i>Videoconferencing over IPv6</i>	30
3 SECONDARY WORK ITEMS	31
3.1 NETWORK MONITORING	31
3.1.1 <i>Partner activities</i>	31
3.2 MULTICAST IPV6	33
3.3 WIRELESS ACCESS	34
3.4 MULTIHOMING	35
3.5 IPSEC	36
3.6 FIREWALLS	36
3.7 OTHER ITEMS	37
4 INTEROPERABILITY (TAHI) TESTS.....	38
4.1 THE TESTED SYSTEMS	38
4.2 OVERVIEW OF IMPLEMENTATIONS AND THEIR TEST RESULTS	38
4.2.1 <i>FreeBSD 4.2 and FreeBSD 4.3</i>	38
4.2.2 <i>AIX 4.3.3</i>	39
4.2.3 <i>Solaris 8</i>	40
4.2.4 <i>Linux 2.2.19 and Linux 2.4.3</i>	41
4.3 SUMMARY	42
5 LIAISON WITH EXTERNAL IPV6 PROJECTS	43
6 CONCLUSIONS AND FUTURE WORK	44
7 REFERENCES.....	47

ACKNOWLEDGEMENTS

This document reports on work carried out jointly by representatives of a number of National Research and Education Networks (NRENs). These people include in alphabetical order, but are not limited to:

- Tim Chown (University of Southampton (UoS) & UKERNA, UK)
- David Harmelin (DANTE, UK)
- Joop Joosten (CERN, Switzerland)
- Simon Leinen (SWITCH, Switzerland)
- Ladislav Lhotka (CESNET, CZ)
- János Mohácsi (BME, Hungary)
- Wiktor Procyk (POZNAN, Poland)
- Juergen Rauschenbach (DFN, Germany)
- Yves Schaaf (RESTENA, Luxembourg)
- Christian Schild (JOIN Project, DFN)
- Miguel Angel Sotos (RedIRIS, Spain)
- Bernard Tuy (RENATER, France)
- Stig Venaas (UNINETT, Norway)
- Wilfried Woeber (ACOnet, Austria)

The author would like to thank UKERNA for funding his attendance of and participation in the TF-NGN IPv6 Working Group, and DANTE [DANTE] and TERENA [TERENA] for their assistance in organising and supporting the TF-NGN activities [TF-NGN].

EXECUTIVE SUMMARY

The GÉANT network is the successor to TEN-155, offering connectivity speeds of up to 10Gbit/s for production IPv4 traffic. However, GÉANT is also committed to provide a native IPv6 service in its lifetime (i.e. by 2003/04), and thus early studies and pilots of IPv6 services are of key importance for the GÉANT project as a whole.

This report describes work undertaken within the IPv6 Working Group of the GÉANT Task Force for Next Generation Networks (TF-NGN). Participants in TF-NGN are generally representatives of the National Research and Education Networks (NRENs) that comprise GÉANT.

The IPv6 activity has undertaken several sub-activities. These have been assigned two levels of priority within the working group. This report covers in detail the first set of activities with higher priority, whilst the second set of activities will be reported upon in an addendum that the group will complete in October 2001.

Throughout the period of the tests reported on in this document, an IPv6 testbed network was maintained, offering connectivity (some native, some tunnelled) for up to fifteen participants through a single core router hosted in Amsterdam. The testbed network enabled successful tests of the use of allocated IPv6 addresses, of a private IPv6 DNS hierarchy (using A6, DNAME, and bitstring records and the ip6.arpa inverse domain), and a number of IPv6-enabled applications.

In addition, most of the participants also undertook the deployment of IPv6 networks within their own NRENs. While much of this connectivity was tunnelled, we can estimate that some 200 IPv6-enabled routers have been running in those networks, the largest probably being that of the JOIN project within the DFN in Germany. A number of IPv4-IPv6 transition tools were successfully tested in some NREN IPv6 networks.

Of the secondary work items, we have encouraging results in each area. Most partners have deployed some methods for monitoring the performance and behaviour of their IPv6 networks; this includes new tools such as trout6 developed by the TIPSTER6 project in Hungary. There have been successful site deployments of multicast IPv6 (using FreeBSD with PIM-SM), IPsec, (non-commercial) firewalls, and wireless 802.11b IPv6-only LANs.

The conclusions of the work are listed in Section 6 below. It is clear that the basic building blocks for a production IPv6 service on GÉANT are in place, but it must be recognised that there is still a significant amount of testing and pilot work to be undertaken.

A proposed Fifth Framework project called 6NET (led by Cisco, including DANTE and a large number of NRENs and universities) is currently under negotiation with the European Commission, and may provide one excellent vehicle for this future work to be undertaken. However, it is also very important that the TF-NGN IPv6 WG activity continues, to ensure appropriate studies are made of non-Cisco implementations and that IPv6 piloting work is open to NRENs not participating in 6NET.

1 INTRODUCTION

This document represents reporting on a broad range of IPv6 activities undertaken by the GÉANT Task-Force Next Generation Networks (TF-NGN) IPv6 Working Group. The group meets 4-5 times per year, studying topical network subjects, such as IP multicast, Premium IP and optical networking. The group's activities often lead to the deployment of future services on GÉANT.

Interest in IPv6 continues to grow. GÉANT is committed to deploying a native IPv6 service by the end of its project lifetime. It is thus important for the TF-NGN IPv6 activities to gain early experience of IPv6 deployment issues, in preparation for a production service at some point within the next 1-3 years.

It is quite probable that the initial native IPv6 deployment on GÉANT will see IPv4 and IPv6 running on the same links, between dual-stack routers. In a production environment the deployment will require hardware support for IPv6 routing, a feature imminent but not yet commercially available in high-end carrier-class products from the likes of Cisco, Juniper or Hitachi.

In this report we overview the objectives of the GÉANT IPv6 studies, outline the network deployments of a cross-section of NRENs, describe the results to date of the IPv6 work items, overview IPv6 collaborations, and finally list the future areas of study for the IPv6 work.

The objectives of the GÉANT IPv6 project (or GÉANT Test Programme, GTPv6) include:

- To gain and develop an understanding of the issues involved in deploying IPv6 networks, in terms of areas including physical infrastructure, address allocation, registries, routing and DNS operation.
- To gain experience in operational management of an IPv4/IPv6 backbone.
- To gain insight into the implications of the new IPv6 protocol, and how it will impact the backbone network, NREN backbones and University end sites.
- To deploy and operate an IPv6 backbone network that can interconnect to any GÉANT participant IPv6 network, and that can peer with/offer transit to other world-wide IPv6 networks.
- To encourage NRENs to participate in GTPv6 such that they in turn may offer IPv6 connectivity to their own universities/sites.
- To collaborate with other European IPv6 projects to gain a better mutual understanding of IPv6 deployment issues.

The methodology for driving forward the GTPv6 activity includes:

- Identification of European projects with mutual interest. The GTPv6 web site contains links to such projects, and members of such projects may be invited to attend GÉANT TF-NGN (GTPv6) meetings.
- Assistance for new sites/participants wishing to join GTPv6.

- Identifying potential new partners/networks with which to liaise outside of the European academic backbone. This may include collaborations with US or Japanese networks.
- General tracking of IETF developments.
- Promotion of the GÉANT IPv6 project.
- Ensuring the GTPv6 web presence is regularly maintained.

The official web page of the TF-NGN IPv6 work [GTPv6] is:

<http://www.ipv6.ac.uk/gtpv6/>

It should be noted that a number of GTPv6 participant NRENs are at the time of writing in negotiation with the European Commission regarding a Fifth Framework project called 6NET, led by Cisco. If funded, 6NET will, as of around January 2002, provide a focus for the deployment of a high-speed native IPv6 network between participant NRENs (including UKERNA, RENATER, DFN, GARR, SWITCH, AConet, SURFnet and NORDUnet). While this should be welcomed as an excellent initiative for pan-European IPv6 academic/research network deployment, there will be other NRENs and router vendors, and experiments outside the scope of 6NET, that require study. Such additional activities will remain under the TF-NGN IPv6 umbrella.

2 PRIMARY WORK ITEMS

The IPv6 activity has undertaken several sub-activities. These have been assigned two levels of priority within the working group. This report will describe in detail the first set of activities with higher priority, whilst the second set of activities will be reported upon in an addendum that the group may complete in October 2001.

Each work item was assigned a leader responsible for coordinating and reporting work done within that area.

2.1 Platforms, routing and interoperability

The TF-TANT [TF-TANT] reporting featured interoperability as one of its four IPv6 work items. At that stage, basic BGP and other router interoperability problems were more commonplace. Such issues have become rarer of late, but are still important to consider (for example between Cisco, Juniper, Hitachi and other ready or near-ready commercial IPv6 routers).

TF-TANT did not undertake formal interoperability tests. In Section 4 below, we detail the results of TAHI conformance and interoperability tests as carried out by BME (Hungary).

2.1.1 Current GTPv6 core IPv6 Network

The GTPv6 core network remains based around a central Ericsson Telebit TBC2000 router, consisting of a TBC2000/3 with a TBC118 controller and a TBC120 controller. The router is also equipped with an ATM switch backplane. The TBC118 provides a 155 Mbps ATM multimode port. The TBC120 controller offers four 10/100 Mbps Ethernet LAN ports and also one 155 Mbps ATM multimode port.

The details of the router configuration are listed on the project web site [GTPv6].

The same router was used in the QTPv6 tests under the TF-TANT work within QUANTUM. As GÉANT comes into service, the TF-NGN IPv6 network should also evolve. It is currently using the same topology and technology as QTPv6, but this is planned to change by the end of October 2001.

The characteristics and limitations of the current IPv6 test bed network are:

- The network is centred on a core ATM router provided by Ericsson Telebit. Thus the topology is a "star" with no sites having redundant/resilient links to other sites. While a good test environment, such a topology is not typical in deployed backbone networks.
- Participants use whatever IPv6-capable router they can spare for their experiments. This typically means that the routers used are either low specification or PC-based routers. The most common platform used is Cisco [CISCO], although until recently (June 2001) Cisco provided only "beta" (not commercially-supported) IPv6 images.
- Participants are able to connect to the core router either via "native" IPv6 links carried in ATM PVCs or by IPv6-in-IPv4 tunnels. The bandwidth of the PVCs is typically 512Kbit/s. While native connectivity is desirable, the bandwidth available for PVCs is a limiting factor.

- Address space is allocated centrally from a 6bone prefix (registered as QTPVSIX) assigned to the GTPv6 project, i.e. 3ffe:8030::/28. Each site (or NRN) that wishes to use GTPv6 address space receives a /34 prefix under that /34 from which it, and its own participant sites, allocated IPv6 addresses.
- BGP4+ was used as the common routing protocol between networks, with the core router on AS8933.
- Very few "real" userland applications have been run on QTPv6 and subsequently GTPv6. The JOIN project reported IRC was the most popular application over IPv6 on their network. Services such as IPv6 DNS have been run with success however.

In summary, the QTPv6 network allowed sites to gain a good understanding of the operation of IPv6-enabled routers, but the resultant network was not of "production" quality, and is not one that in its current form could be utilised as such. Most participants had other external IPv6 links and peerings through which they could also exchange IPv6 traffic.

GTPv6 Network, July 2001

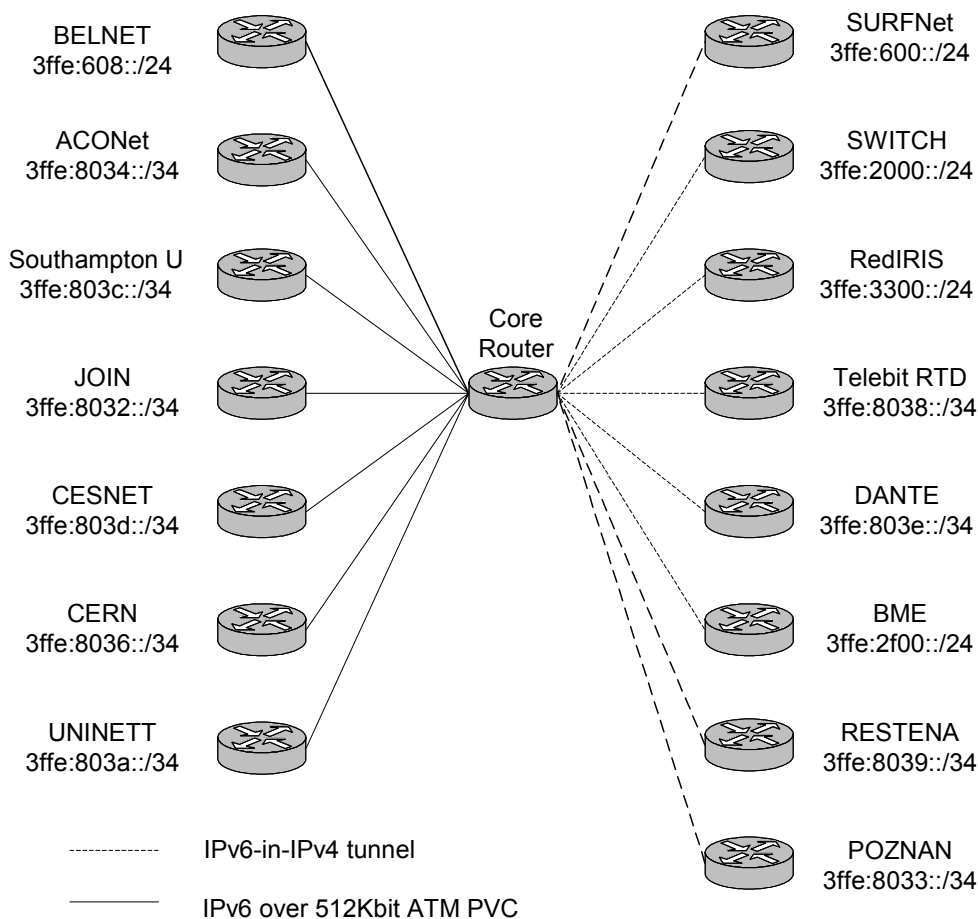


Figure 2.1: GTPv6 participant connectivity to the core router

Address space has been allocated to participants (where requested) as listed below. The 6bone address allocation boundaries are slightly different from the production prefixes (where SubTLAs are /35's, growing up to /29's). The address allocations were made using the

“flexible method for managing the assignment of bits of an IPv6 address block” IETF I-D [BLANCHET].

Bits	28	6	14	16	64
Field	pTLA	NLA	SLA	SubSLA	Interface ID

Prefix	Holder
3ffe:8038:0000::/34	QTPV6 core network
3ffe:8034:0000::/34	ACOnet
3ffe:803c:0000::/34	Univ. of Southampton / UKERNA
3ffe:8032:0000::/34	JOIN
3ffe:8036:0000::/34	CERN
3ffe:803a:0000::/34	UNINETT
3ffe:803e:0000::/34	DANTE
3ffe:8031:0000::/34	-- available --
3ffe:8039:0000::/34	RESTENA
3ffe:8035:0000::/34	ARNES
3ffe:803d:0000::/34	CESNET
3ffe:8033:0000::/34	POZNAN
3ffe:8037:0000::/34	HEAnet

Table 2.1: GTPv6 participant address allocations within 3ffe:8030::/28

Each NREN is free to adopt its own assignment policies within its national network, e.g. CESNET allocates a /42 to regional network centres, and a /48 to each end sites (see below). The common registry policy [ALLOC] as agreed by ARIN, APNIC and RIPE is that all IPv6 sites should be allocated a /48 network.

2.1.2 Evolving the GTPv6 core network

The GTPv6 group feels that it should seek to deploy well-specified IPv6-purposed routers at the earliest possible opportunity. Emphasis should be placed on establishing a more production-like quality (or as close as possible given funding constraints) network for experimentation.

It is quite possible that such a network may emerge from the 6NET project. However we must consider trials of other router vendors' IPv6 equipment and testbed access for non-6NET NRENs. It is also not certain that 6NET will pass through negotiation to completion.

There are two constraints on the router interconnections:

1. The new GÉANT network will not be ATM-based; thus there will not be the ability to run native PVC-based connections beyond the date at which such PVCs will be removed, i.e. 1st December 2001. It should also be noted that many NRNs will discontinue ATM use well before that date. However, some NRENs (e.g. RENATER) have kept ATM equipment with a view to supporting ongoing PVC-based connectivity.
2. There are no fibre circuits that could be utilised for dedicated native IPv6 interconnections within the GÉANT backbone infrastructure at the present time (this would be expensive, as would be provision of dedicated local loop links).

DANTE has experimented successfully with MPLS encapsulation on Juniper routers using Circuit Cross-Connect (CCC) for link-layer forwarding. Detailed results are subject to NDA, but tests have proven successful at speeds up to 622Mbit/s. Future use of MPLS on GÉANT is not yet clear, but it may be introduced for other purposes (e.g. QoS provisioning) and thus also be available for IPv6 encapsulation (if such encapsulation can provide the necessary IPv6 testbed properties).

There are four main scenarios for native IPv6 adoption on the GÉANT core backbone and routers:

1. Running native IPv6 as a secondary protocol to IPv4 on dual-stack production GÉANT backbone routers, with IPv4 and IPv6 running on the same links. Gradually more and more traffic becomes IPv6. At present this is not the more likely option since most router vendors have no IPv6 carrier-class routing product with hardware support.
2. Running IPv6 on parallel fibre infrastructure. Such an IPv6-only network would eventually become prominent. This is clearly the more costly option in terms of resources. It also has a disadvantage if the native bandwidth is significantly less than that available over IPv6-in-IPv4 tunnels.
3. Running IPv6 tunnelled in IPv4, via either manually configured tunnels or running 6to4, leveraging the high-performance IPv4 routing infrastructure. As and when the infrastructure allows native IPv6 links, or dual stack operation as per option #1 above, tunnels are removed and replaced by such native links. Realistically, this is the most probable option.
4. Running IPv6 encapsulated at the link layer using techniques such as Juniper's CCC for MPLS. Cisco has an equivalent method called AToM [AToM], and has also recently made available a more efficient encapsulation technique (that does not use MPLS) called Universal Tunnel Interface (UTI).

The GTPv6 group feels that for experimentation deployment of IPv6-purposed routers should take priority over the requirement for native IPv6 connectivity, i.e. tunnelled IPv6-in-IPv4 connections would, initially, be acceptable.

The GTPv6 IPv6 network in the short-to-medium term is likely be characterised by:

- A topology more representative of a "real" backbone network, i.e. core routers with a number of multiple interconnections. This is in contrast to the QTPv6 "star" yet would most likely not extend to a full mesh. It is unlikely that there will be a single "core" router.
- Participants will ideally use well-specified router hardware from commercial vendors where possible. We would for example seek to deploy Juniper and Hitachi IPv6 routers on the GTPv6 network. However, PC-based routers, such as FreeBSD, may offer more flexibility due to their open-source nature (e.g. in working on inter-domain IPv6 multicast with PIM-SM).
- Participant NRENs will be encouraged to utilise SubTLA address space where possible, i.e. production IPv6 address space obtained from RIPE under the 2001::/16 prefix. This stresses the importance of address allocation policy issues. Sites at very early stages of investigating IPv6 may be best advised by their NREN to connect via 6bone prefixes/tunnels.

- BGP4+ will likely remain the routing protocol between networks.
- The adoption, porting and development of user applications will be encouraged, though it is appreciated that GTPv6 participants may lack the time to port or develop new software. End sites within NRENs may have the resources, especially on European funded projects. Vendor-supplied software may be expected also.

The GTPv6 ATM PVCs will no longer be available when TEN-155 ends on the 1st December 2001.

2.1.3 IPv6 Support in Routers

There are a number of router products with IPv6 support. These include:

- Cisco IOS, as of version 12.2(2)T (commercial support) [CISCO]
- Ericsson Telebit, e.g. on the TBC2000 used on GTPv6, commercially supported
- GateD
- Hitachi GR2000 range
- Juniper (JUNOS), announced IPv6 product for Japan in 4th Quarter 2001, early access is likely to be available for GTPv6 via a Juniper M5 router to be hosted by RENATER.
- 3Com NetBuilder range
- Zebra (www.zebra.org), running on at least *BSD and Linux.
- Multi-Threaded Router (MRT)
- 6WIND IP Edge Device

BGP4+ is almost universally available for exterior routing, but availability of interior routing protocols for IPv6 (OSPF, RIP, IS-IS) is currently mixed.

However, as yet none of the vendors have hardware support for IPv6; thus running a dual-stack router will have a performance impact on the IPv4 traffic as router processor cycles are used to handle IPv6 traffic.

A complete router product list is maintained at the IPv6 Forum [IPV6FORUM] web site.

2.1.4 UNINETT (Norway)

UNINETT is offering IPv6 access to all Norwegian universities, a few colleges and some research institutions. Some have native access and some have tunnels. The 6to4 method is used for dynamic IPv6-in-IPv4 tunnelling. The universities are doing IPv6 work on their own and testing a variety of software, sometimes in cooperation with UNINETT.

UNINETT runs native IPv6 STM-1 links, some ATM, some with native IPv4 and IPv6 in parallel, and some IPv6-in-IPv4 tunnels. There are several BGP peerings established to external IPv6 networks. Internally, RIP is used for routing IPv6. See Figure 2.2 for the topology.

UNINETT has been allocated 2001:700::/35 from RIPE, and also has two 6bone prefixes. Production addresses have been deployed in the backbone and delegated to some customers.

Further UNINETT IPv6 information: <http://www.uninett.no/testnett/index.en.html>

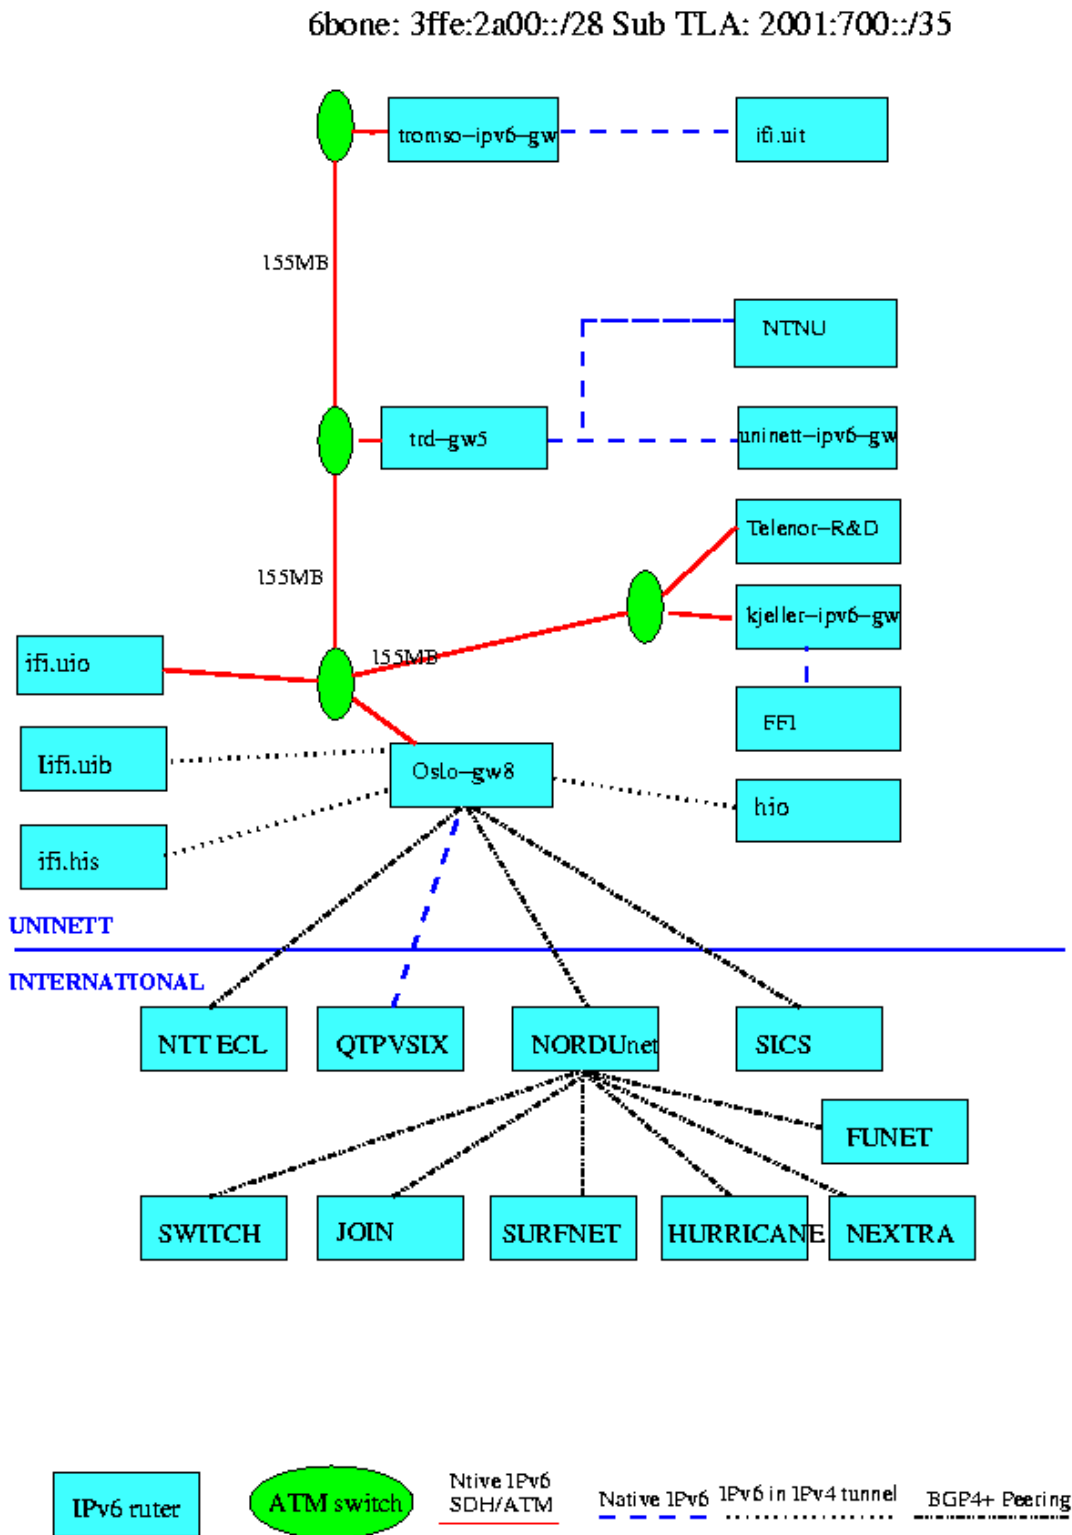


Figure 2.2: UNINETT IPv6 Network Deployment

2.1.5 UKERNA (UK)

UKERNA is managing a one-year JISC-funded IPv6 deployment study for UK academic sites called Bermuda 2, the participants being University of Southampton (UoS), Lancaster University and University College London (UCL).

The original method for the interconnectivity of the Bermuda 2 partners was a set of three ATM PVCs running over the SuperJANET III managed bandwidth service (MBS).



Figure 2.3: UK IPv6 deployment: Bermuda 2

For the initial part of the project, the native ATM PVCs were available and running. However, the demise of ATM (when SuperJANET III was replaced by SuperJANET 4), and the removal of ATM by many of the UK regional academic MANs, has meant that much of the PVC provision has been lost. At the time of writing UCL connects to Lancaster via an IPv6-in-IPv4 tunnel, while UoS still connects to UCL via a native PVC link. The link bandwidth on the ATM PVC is 2Mbit/s; the tunnelled link runs at the IPv4 connectivity bandwidth, possibly up to Gigabit speeds on the SuperJANET backbone.

There is connectivity from UCL via an ATM PVC to BT's test bed network; this link was used by UoS and BT on the Fifth Framework project 6INIT [6INIT].

The Bermuda project routers at UCL and Lancaster are Cisco platform, while UoS uses an Ericsson Telebit TBC2000.

UCL has a 20Mbit/s native IPv6 path direct to Japan which is currently the subject of collaborative talks between EU partners (led by RENATER, but including TF-NGN IPv6 WG participants) and WIDE in Japan.

The UCL-CRL link provided network connectivity from the IPv6 Forum Summit, via the TEN-155 network, to Beppu and CRL for a multi-way IPv6 conferencing event in January 2001 (see Figure 2.4). The link is available until at least November 2001.

Further Bermuda 2 information is available at: <http://www.ipv6.ac.uk/bermuda2/>

ATM Configuration

19 Jan. 2001 CRL

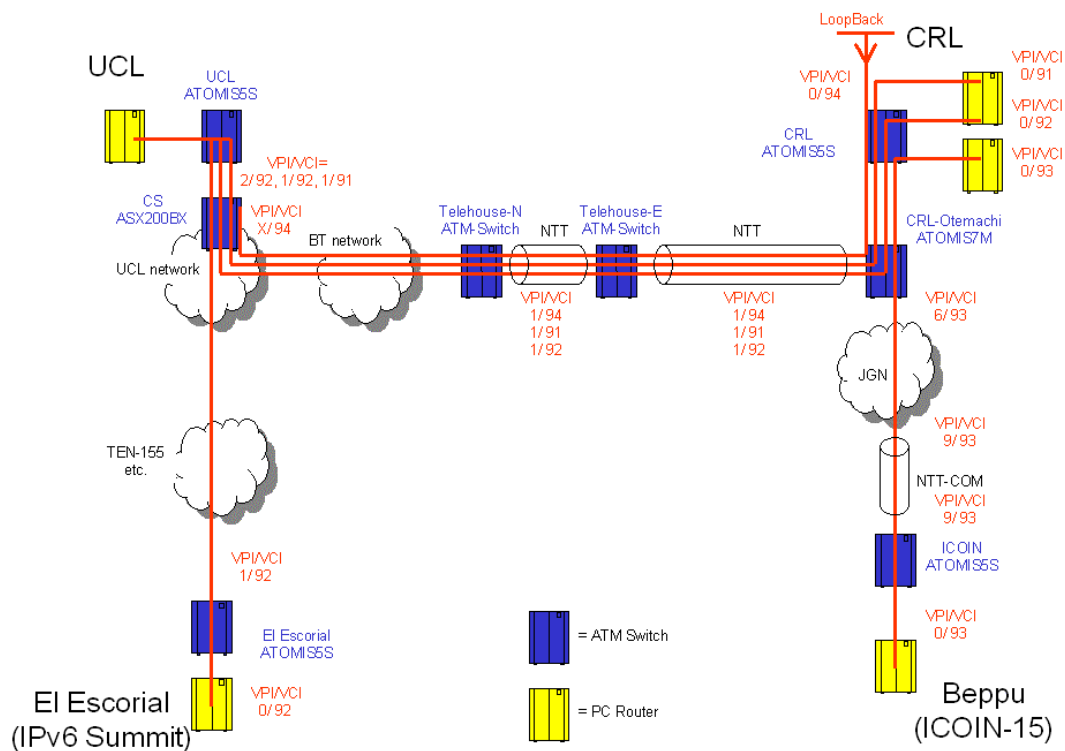


Figure 2.4: IPv6 connectivity from El Escorial to Beppu, via TEN-155, as of January 2001

2.1.6 RedIRIS (Spain)

The RedIRIS IPv6 network in Spain is currently based around a central router from which IPv6-in-IPv4 tunnels run out to the autonomous regional networks.

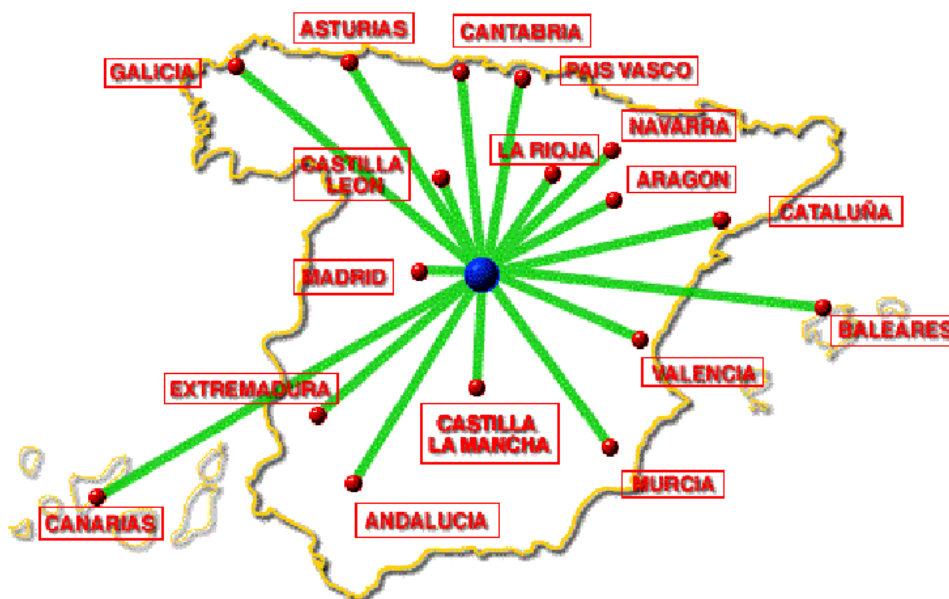


Figure 2.5: RedIRIS IPv6 network topology

Approximately 70% of the network is IPv6-connected, and the network has been stable for basic IPv6 functionality since January 1998.

RedIRIS is allocating address space from its RIPE prefix (2001:0720::/35) and also configuring reverse address lookup delegation. It also uses its 6bone prefix (3ffe:3300::/24). A new IPv4 network is being designed, but the nature of the equipment is as yet not known, thus it is not possible to say for sure what the IPv6 connectivity method will be.

The current ATM network will soon be used for IPv6, with initially low bandwidth PVCs running to each region, then later it is likely that the links will be upgraded to 155Mbps ATM or SONET.

With the 6bone addressing, each region gets a /32 network, with sites getting a /48 allocation each. Routing platforms used are Cisco (IOS 12.0(T)), Ericsson Telebit and Linux.

For IPv6 site connectivity, each site connects via an IPv6-in-IPv4 tunnel to the nearest (in the IPv4 sense) IPv6 access router.




Figure 2.6: RedIRIS Access Nodes

RedIRIS has international IPv6 connectivity tunnels to Mexico, Chile, GTPv6 and SURFnet.

For further RedIRIS IPv6 information, see: <http://www.rediris.es/red/iris-ipv6/>

2.1.7 JOIN / DFN (Germany)

JOIN has offered IPv6 connectivity to the 6bone for its users for several years. The central 6bone backbone router is located in JOIN's labs in the University of Muenster and connects more than 110 German and about 10 international leaf sites. The backbone has around 20 international IPv6 connections/peerings (see Figure 2.7, though this is a little out of date).

An up to date list of all JOIN's connected IPv6 leaf sites can be found at: <http://www.join.uni-muenster.de/6bone/6bone-join-nla-e.html> .

This service is not limited to DFN customers. Only a third of the leaf sites are research and educational facilities. The rest consist of commercial development companies, private individuals, schools and student residences.

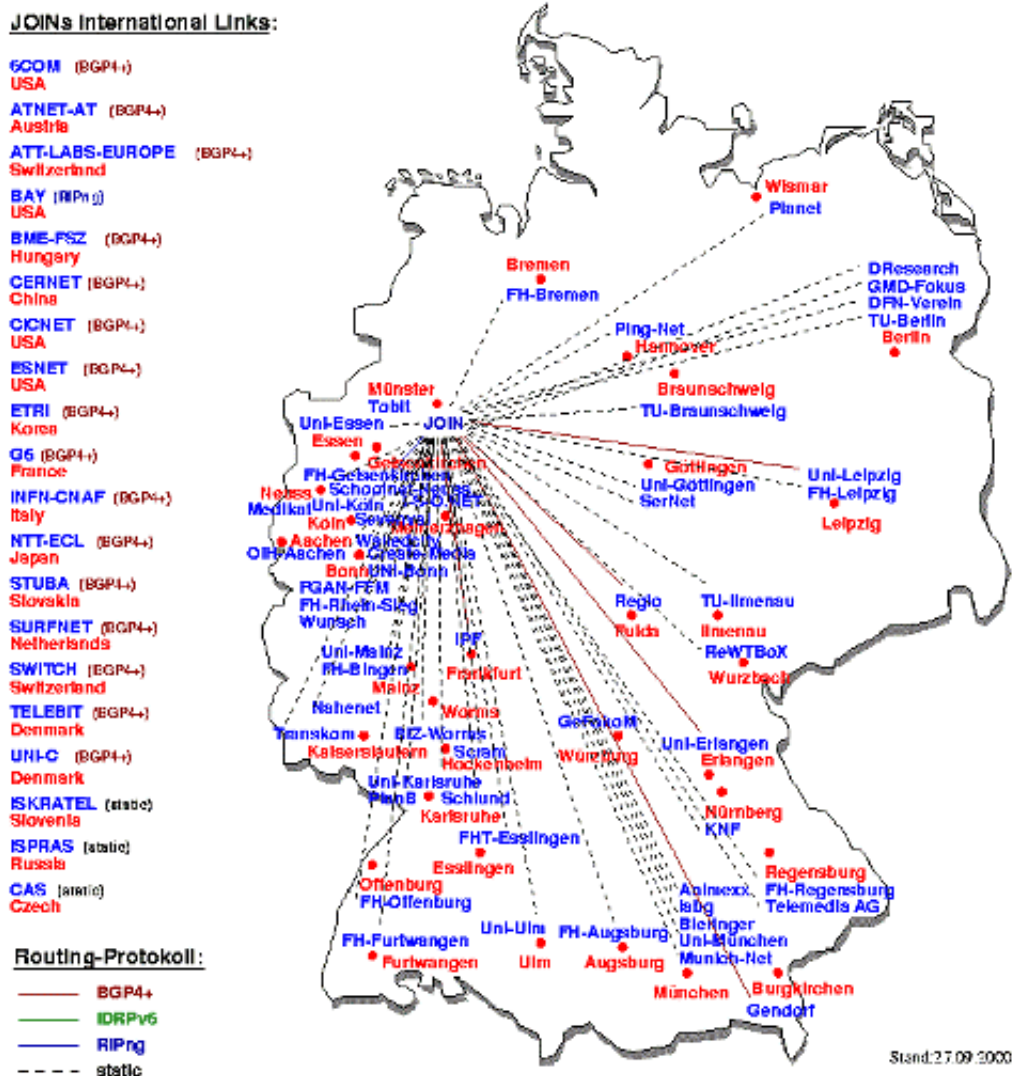


Figure 2.7: JOIN Project leaf sites (as of September 2000)

DFN recently started to deploy prefixes from their own production address space (2001:638::/35) to move away from the test 6bone network to a more production network. Some of DFN's customers changed their 6bone prefix to DFN's RIPE-allocated prefix. Increasing vendor support for IPv6 has led to many new customers getting activated and connected, using the new RIPE prefix.

As part of the transition, JOIN separated the tunnel services from the older 6bone backbone router and connected all DFN customers through a second dedicated router (as illustrated in Figure 2.8).

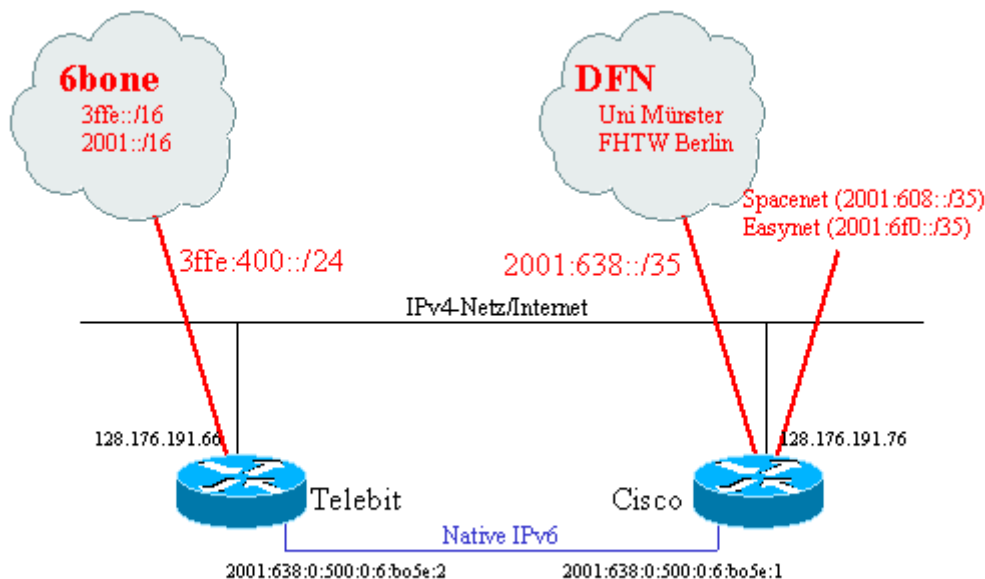
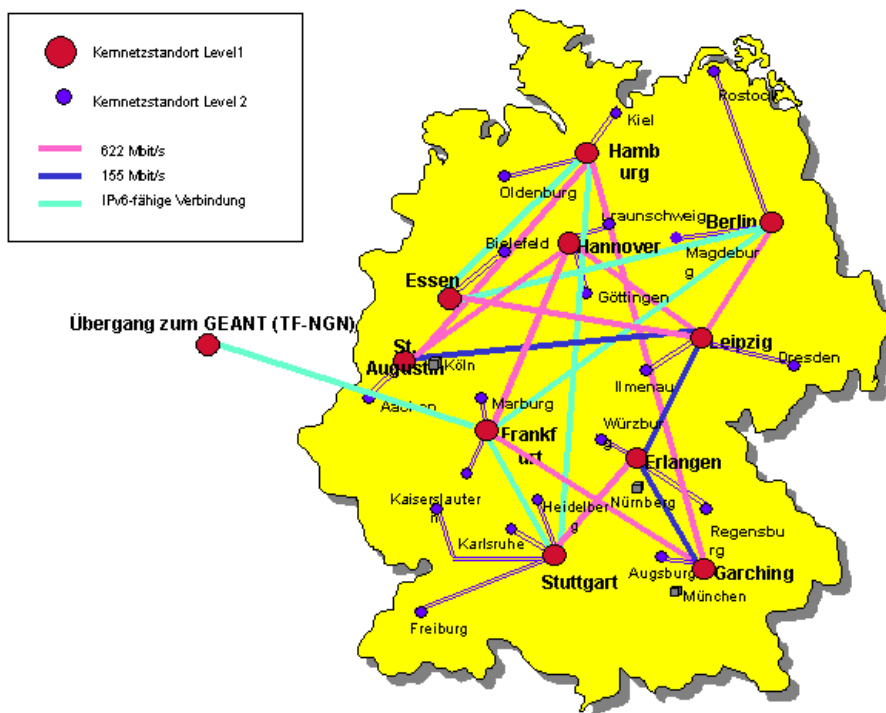


Figure 2.8: JOIN's split service routers

So far four new institutions are connected. Two connections go to other providers that are directly connected to DFN's network and only have a RIPE production prefix.

Gigabit-Wissenschaftsnetz (G-WiN)



eitner\leigene Dateien\ppt\quinkart\gw_kernetz_dtsch_komplett.ppt

Figure 2.9: Candidate topology for next phase JOIN IPv6 network

This centralized Cisco tunnel server is the first step in the integration of an IPv6 backbone in the DFN network. JOIN is planning to build an IPv6 network consisting of five dedicated IPv6 routers connected over dedicated fibres (STM1). These five routers will be spread all

over Germany, though the exact location is not yet decided (a candidate topology is shown in Figure 8). These installations will be done at the end of this year (2001).

JOIN and DFN offer some IPv6-ready services to the public and DFN's own customers. These include DNS, ftp, 6to4 and a tunnel broker.

For more information on JOIN, see: <http://www.join.uni-muenster.de/welcome-e.html>

2.1.8 CESNET (Czech Republic)

The Czech National Research and Education Network operator (CESNET) has an IPv6 deployment that, as with many other national deployments, offers tunnelled connections from a single central router connected to the GTPv6 core router.




Figure 2.10: CESNET IPv6 network

The CESNET network allocates a /42 (from the GTPv6 /34 prefix) to each of the six network regions. It also uses the production RIPE 2001:0718::/35 prefix. There are currently twelve customer end sites, each with a /48 site network assigned to them.

For more information, see: <http://www.ten.cz/english/project/ipv6/>

2.1.9 POZNAN (Poland)

No network connectivity map is currently available for the emerging Polish academic IPv6 network. However, POZNAN does have a set of objectives through which it aims to introduce an IPv6 network service in the coming few months.

The objectives of the implementation of IPv6 in the Polish NREN POL-34 network are:

1. To deploy and operate a production quality native IPv6 network connecting particular Metropolitan Area Networks within the Polish network:

First phase:

- Deployment of a production quality local IPv6 network in PSNC
- Deployment of a national IPv6 network

- Establishing a connection to 6bone.pl in Poland: <http://www.6bone.pl>. This will be based on software routers (Linux or/and FreeBSD and Zebra routing software). The connection will be based on ATM PVC channels and/or IPv6 over IPv4 tunnels
- Deployment of basic IPv6 services including ftp and www

Second phase:

- Deployment of native IPv6 connections based on optical network technology (one λ will be dedicated to IPv6 traffic) and dedicated hardware routers.
2. To make international connections to operational IPv6 networks and testbeds in the European Union and Eastern Europe and to the 6bone network. Planned external connections include:
 - A connection to the GTPv6 experiment (being established at the time this document was authored).
 - A connection to the 6bone.
 3. To deploy and use a number of applications supporting IPv6 such as: www, ftp, telnet, ssh, sendmail, news, DNS, etc.
 4. To manage the IPv6 network
 - To experiment with porting existing open source SNMP applications from IPv4 to IPv6
 - To develop management applications for IPv6
 - To establish and maintain VPNs using IPv6
 - To develop tools for management of IPv6 Virtual Private Networks

The Polish experiments within POZNAN have led to some initial findings, which will be reported on in some detail in the Addendum to this deliverable. The current status of the experiments is:

1. Accomplished tasks
 - Tests of various operating systems supporting IPv6: Microsoft Windows NT 4.0/2000 + patch, Linux (Mandrake, SuSE, RedHat, PLD), FreeBSD, NetBSD, OpenBSD.
 - Tests of PC-based platforms for deployment software routers integrated with ATM networks.
 - Tests of server and client applications: www, ftp, ssh, telnet, etc.
2. Tasks in progress
 - Hardware router interoperability tests
 - Configuration of dynamic routing
 - Setting up external connections to GTPv6 and the 6bone.
3. The main activities and objectives for the near future are:
 - To work out the recommended software router configuration supporting both IPv6 and ATM interfaces
 - To deploy a native IPv6 operational network in Poland based on ATM technology (POL-34)
 - To investigate, experiment with and pilot the current state of the art of management applications with IPv6 support.

Note that other NRENs are deploying IPv6 pilot networks; those reported here are a representative sample of all the work.

2.2 DNS

One of the most critical network services for IPv6 (and IPv4) is the Domain Name Service (DNS). The DNS maps domain names to IP addresses, and vice versa. In IPv4, DNS A records are used to indicate IPv4 addresses. In the early days of IPv6 design (circa 1995), a simple method was proposed to represent IPv6 addresses in the DNS. Rather than plain A records, an AAAA (or “quad A”) record is used to tie an IPv6 address to a name [AAAA]. For inverse lookups (IP addresses to names), a new ip6.int domain was devised.

Deployment of AAAA and ip6.int is commonplace for IPv6 networks today, though the number of IPv6 networks is still very small compared to the IPv4 Internet. However, in July 2000 an alternative IPv6 DNS scheme moved to RFC status, namely the A6 system [A6]. Under A6, it is possible to have indirection in the DNS lookups, making it possible to delegate DNS maintenance by network topology and, thus in theory, to better facilitate IPv6 network renumbering. New DNS features, DNAME and the bit-string label, assist the prefix delegation process. Also, the inverse lookup domain becomes ip6.arpa.

The aim of A6 is to facilitate (rapid) network renumbering, making DNS maintenance simpler for systems administrators. However, with AAAA already deployed, transition to A6 would require client systems to understand and act upon the new A6 record types.

Under the TF-TANT [TF-TANT] work, initial investigations were made of the use of AAAA and A6 records for IPv6 DNS entries, where A6 records were tested under an early beta release of BIND9. Since then, BIND9 has become released code, and has support for the new resource record types (A6 and DNAME) as well as DNS lookups over IPv6 transport.

The choice of use of AAAA or A6 records is the subject of considerable debate at present, e.g. [AUSTEIN, BERNSTEIN]. Further, the 51st IETF meeting in London (August 2001) recommended in a joint ngtrans/DNSEXEXT session that AAAA be the deployment of choice, with A6 being moved to Experimental status. If this decision is approved formally by the IETF, though it is not certain that it will, then developers will in future support AAAA and probably drop A6, since Experimental status implies that A6 should not be used in a production environment.

However, one can argue that there is still an interest in gaining operational experience of A6. One of the claimed advantages of A6/DNAME is that they make renumbering easier, but it is not clear in the future Internet how often renumbering will be required, and to what scale (there are also other technical hurdles to jump for site renumbering). There is potentially good reason to use A6/DNAME if (rapid) renumbering is required, but at the same time one could argue for using A6 0 (no indirection, thus similar to AAAA) for simplicity where possible.

While ip6.int is deployed and delegated across the 6bone, ip6.arpa is not; this is causing problems for IPv6 deployment in general. The GTPv6 test DNS servers are using ip6.arpa. It is not clear whether the inverse delegations will remain under ip6.int if A6 becomes Experimental, though that would be the most likely outcome.

It should be noted of course that the A6 tests run by the GTPv6 participants were undertaken before the debate at the 51st IETF meeting.

2.2.1 DNS activity by partners

DANTE operates an ip6.arpa. root server at ws2.nl.ten-155.net.

JOIN operates a DNS server running the latest production releases of BIND-9. With this server JOIN offers not only standard AAAA and IPv6 reverse nameservice, but also offers a test environment for A6 and DNAME records for all interested parties, including A6/DNAME chains. The nameserver creates its own ip6.arpa. root domain if necessary, but usually customers are supposed to use DANTE's root server. That nameserver starts a chain for A6 records that is usable for test purposes for every one of JOIN's customers.

UNINETT is testing BIND9, A6, and DNAME, but using BIND 8 and AAAA and reverse entries for IPv6 on production DNS servers. It has released a patch for the BIND 8.2.3 resolver that makes it use IPv6 transport. You can then put IPv6 addresses for nameservers in resolv.conf and link programs with the 8.2.3 resolver (-lbind) rather than your current resolver libs (see <http://www.venaas.no/dns/bind/> for the patch). Note that at least BSDs with KAME [KAME] and Linux with glibc 2.1.92 or newer (including RedHat 7) already have resolvers that can do this. One could also use the lightweight resolver that comes with BIND 9, or run BIND9 locally and let the current resolver use IPv4 to localhost.

SWITCH has been using the new A6 records and binary labels/DNAME delegation and kept a DNS name server installation up to date. An existing script that is used to generate inverse mappings from "forward" zone files was modified to convert AAAA records into both types of IPv6 inverse records in use.

RedIRIS is running BIND 8 and BIND 9, and is working with ES-NIC to start registering .es for IPv6 addresses and to build a Spanish DNS hierarchy for IPv4 and IPv6, initially a different hierarchy, then merged to be the same.

2.2.2 Test results

Tests by DANTE have shown that BIND 9.2 can synthesize AAAA records from A6 records. The test root server at ws2.nl.ten-155.net was upgraded to BIND 9.2.0a2, and the configuration was modified to include:

```
allow-v6-synthesis { any; };
```

This was tested by querying an AAAA name, which was defined as an A6 chain (2 elements), and it successfully returned the AAAA result. It can also synthesise PTR responses to nibble-based inverse queries under IP6.INT from the bit-string-based PTRs under IP6.ARPA (that are used in GTPv6).

So it seems that it would not be that difficult to have the old-style inverse domain for the GTPv6 prefix delegated to the working group (and thus make traceroute results through our network more human-friendly for people running network tests).

GTPv6 could maintain the inverse domains under .ip6.arpa; if 6bone want to delegate under ip6.int, they could do it like this:

```
0.3.0.8.e.f.f.3.ip6.int      86400 IN   NS   truc.dante.org.uk
0.3.0.8.e.f.f.3.ip6.int      86400 IN   NS   truc.switch.ch
```

Then we would need to ensure truc.dante.org.uk and truc.switch.ch resolve AAAA, A (and A6) requests, and that our BINDs reply to requests under ip6.int (which we have successfully demonstrated possible already).

Some reverse lookups are set up on ws2.nl.ten-155.net (dns.ipv5.ten-155.net), under ip6.arpa:

```
; Reverse DNS for the qtpvsix core is in the ipv6.ten-155.net
zone
\[x3ffe80380080/48]      DNAME    ipv6.ten-155.net.

; Delegation to aconet
\[x3ffe80340/34]      DNAME    v6.aco.net.

; Delegation to switch
\[x3ffe20/24]         NS       ns1.ipv6.switch.ch.
                       NS       ns2.ipv6.switch.ch.
\[x20010620/29]      NS       ns1.ipv6.switch.ch.
                       NS       ns2.ipv6.switch.ch.
\[x2002823b01d2/48]  NS       ns1.ipv6.switch.ch.
                       NS       ns2.ipv6.switch.ch.

;Delegation to dfn
\[x3ffe0832/32]      DNAME    ip6.qtpsix.join.uni-
muenster.de.
\[x3ffe04/24]        DNAME    ip6.6bone.join.uni-
muenster.de.
\[x20010638/29]      DNAME    ip6.dfn.join.uni-muenster.de.
```

The immediate next step is to secure the ip6.int delegation and to use the synthesis method to have the existing .ip6.arpa hierarchy return reverse lookups for traceroutes across GTPv6.

Clearly A6 could be deployed, and AAAA records could be synthesised from A6, but the issue of A6 complexity is a genuine concern, and we also lack experience in testing AAAA synthesis for non-trivial cases.

In the longer term, the GTPv6 group will need to decide whether to continue with A6 trials; this is expected to be discussed at the next TF-NGN meeting in Athens in October 2001. GTPv6 does not operate a production network, but it is debatable whether A6 should be studied if its chances of real deployment in the near to medium future are slim. We should know the formal IETF decision by October.

2.3 Registries and Addressing

There are two “classes” of IPv6 address being allocated at present. The older 6bone [6BONE] addresses are allocated from the IPv6 test prefix of 3ffe::/16. Prior to July 1999, when the regional registries started allocating production IPv6 addresses, it was only possible to use 6bone addresses.

Although both prefixes are available, and most GTPv6 NRENs have production prefixes, the older 6bone prefixes are also still used. The 6bone addresses can be used for more experimental testbed networks, and some sites may wish to use both types of addresses so that they can perform IPv6 multihoming tests.

Since RIPE and the other registries began allocating production IPv6 SubTTLA addresses, RIPE has been leading the “race” for assignments (20 in ARIN, 31 in APNIC, 41 in RIPE, 92 total) [ADDRALLOC].




Figure 2.11: IPv6 Production Prefix Allocations by region

NRENs with production address space include (see <http://www.ripe.net/cgi-bin/ipv6allocs>):

```
RIPE-NCC (whois.ripe.net)
ES-REDIRIS-20010521      2001:0720::/35
CZ-TEN-34-20010521     2001:0718::/35
NO-UNINETT-20010406    2001:0700::/35
BE-BELNET-20001101     2001:06A8::/35
FR-RENATER-20000321    2001:0660::/35
GR-GRNET-19991208     2001:0648::/35
DE-DFN-19991102       2001:0638::/35
UK-JANET-19991019     2001:0630::/35
AT-ACONET-19990920    2001:0628::/35
CH-SWITCH-19990903    2001:0620::/35
NL-SURFNET-19990819   2001:0610::/35
```

RIPE has information on IPv6 at: <http://www.ripe.net/ipv6>. It appears that the registries will be prolonging the initial 100 SubTLA bootstrap phase; this means it will continue to be relatively easier for top-level providers to get production IPv6 address space.

The Provisional IPv6 Assignment and Address Allocation Policy Document [ALLOC] agreed upon by all the registries suggests the adoption of /48 allocations for sites. However the definition of a site, e.g. a home network or a whole university, is not well-defined. There is certainly a feeling among the GTPv6 participants that the /35 SubTLA slow roll-out to NRENs should be lifted, such that NRENs can begin address plans from a full /29 prefix (see also a RIPE IPv6 WG review of the Policy Document [196REV], which also suggests making the qualification criteria for assignment more relaxed). In fact, it may be desirable to move to a SubTLA boundary at /28 or /24 (as per the 6bone) to make the byte/nibble alignment cleaner.

JOIN has allocated addresses to a large number (over 50) leaf sites under their 6bone prefix using RFC2471 (which specifies an IPv6 Testing Address Allocation scheme).

The GTPv6 core addresses are assigned via “A flexible method for managing the assignment of bits of an IPv6 address block” [BLANCHET].

2.4 Transition Tools

IPv6 will not be deployed overnight, nor will there be a “flag day” where we change from IPv4. Instead, mechanisms are required for IPv4 and IPv6 to co-exist, and for IPv6 transition to occur gracefully.

There are three levels at which the transition will occur:

- *Backbone networks.* In the case of European NRENs, the backbone network is GÉANT. GÉANT is committed to introduce IPv6 in its lifetime. However, it will also need to carry production IPv4 traffic. Methods for the transition were outlined in Section 2.1.2; a likely scenario is that the core routers will support IPv4 and IPv6 in hardware, dual-stack, and native IPv4 and IPv6 will co-exist on the same links. There may also be use of tunnelling and (MPLS) encapsulation methods.
- *Within the NRENs.* Each NREN may have different network topologies and technology. Some still use ATM, and can thus run IPv6 over ATM PVCs, others may not, but may consider IPv6 encapsulation in MPLS. Alternatively manual or 6to4 tunnelling can be used to provide IPv6-in-IPv4 connectivity. Provision of a tunnel broker allows individual users or networks to connect to the IPv6 network, provided IPv4 firewalls don't block Protocol 41 (IP-in-IP).
- *At end sites (universities).* Here, each site needs to provision IPv6. In an academic environment, IPv4 addresses are likely to be in good supply, so running an entire site dual-stack is a possibility. However, if IPv6-only networking is desired, or IPv4 addresses are scarce, methods such as DSTM, NAT-PT and dual-stack application-layer proxies (e.g. SMTP relays, web caches, irc servers) can offer IPv6 functionality.

An example of one NREN's considerations on transition can be found in the document written by Stig Venaas of UNINETT [UNI2000]. This is an excellent overview of the technologies, and the reasonable steps an NREN might take in transition.

2.4.1 Partner activities

UNINETT has developed an experimental tunnel broker ([at tunnelbroker.uninett.no](http://tunnelbroker.uninett.no)) which has approximately 600 users registered to date (from all over the world – this highlights a problem in that the tunnels may not use topologically efficient IPv4 paths for the IPv6 tunnels).

UNINETT is also running a 6to4 gateway for internal use, and a 6to4 relay for public use. It may want to restrict access to the relay somewhat or try to find ways to monitor the use. However, a 6to4 relay for GTPv6 NRENs is a desirable facility.

SWITCH is using IPv6-over-IPv4 tunnels to connect to other parts of the 6bone and to connect new sites to its part of the 6bone. It has three 6to4 gateways (which should be consolidated), and one of them is now set up with the RFC3068 anycast address. This has been used successfully as the default gateway on a home router, which has a 6to4 tunnel interface. The anycast method is required for 6to4 gateways to find a default route.

Within its own site, SWITCH runs IPv4 and IPv6 in parallel on a few Ethernet and serial links. All workstations that have IPv6-capable Oses (Linux or Solaris 8) and that live on segments with IPv6-capable routers are configured to use IPv6. They usually get two or three global addresses; one from the SWITCH 6bone prefix (3ffe:2000::/24), one from its Sub-TLA (3ffe:0620::/35) and one from one of the SWITCH 6to4 blocks (2002:823b:::/48). Most hosts have their IPv6 addresses listed under ipv6.switch.ch rather than switch.ch, so nobody

notices they have IPv6. If both the A and the AAAA records resolve to the same host, you would transparently use IPv6. It would be interesting to try this configuration on machines that provide services such as NFS, print spooling, NIS etc. SWITCH also intends to build a small IPv6-only LAN and to experiment with tools like NAT-PT to access the IPv4 Internet.

JOIN (DFN) offers a 6to4 gateway in Regensburg and a tunnel broker in Leipzig (see <http://joshua.informatik.uni-leipzig.de>).

The interesting IPv6 transition challenge comes from NRENs that have migrated away from ATM; options for native IPv6 are limited until the network routers can run dual-stack. However, some vendors including Cisco, Juniper and Hitachi have announced plans for IPv6 support in hardware within the next 6-9 months. Until then, tunnels (manual or 6to4) are the most natural option.

Within end sites there needs to be more work done studying requirements for IPv6-only sites, which includes studies of when to use NAT-PT against when to use application layer proxies, and detailed analysis of IPv4-IPv6 DNS interactions.

2.5 Applications

Although applications per se aren't required in deploying an IPv6 network across GÉANT and the NREN networks, such applications are critical for proving the network technology.

The development by BME of an IPv6 applications database is thus a very welcome development in tracking IPv6 ports and statuses of a wide variety of software packages.

The applications may run on a variety of operating systems. The following systems all have IPv6 implementations released for them:

- Solaris 8, installed optionally at system build time, commercially supported
- FreeBSD, included in FreeBSD 4.3, also available as a patch from the KAME project
- OpenBSD and NetBSD
- Windows NT, available as a patch
- Windows 2000, available as a "technology preview" patch, requiring Service Pack 1
- Windows XP, will include IPv6 on the distribution CD but not installed by default
- Linux, via various sources including the USAGI project
- AIX 4.3
- Tru64, e.g. for the Alpha platform
- HP/UX 11.0

However, some implementations are more feature rich than others, the *BSD variants being the most advanced.

Under Windows, IPv6 is available as a stack, with development support for C, but, for example, most Windows applications are not built to use the stack, Internet Explorer being one exception.

Solaris 8 is currently the only commercially supported IPv6-enabled operating system.

The IPv6 Forum web site [IPV6FORUM] lists the available operating systems.

2.5.1 TIPSTER6 Applications Database

Several operating system vendors have started to provide their systems with IPv6 support. One of the largest problems facing would-be IPv6 users is the lack of native IPv6 connectivity or the shortage of IPv6 compatible applications. Unfortunately, the legacy IPv4 client applications cannot use the IPv6 services without some kind of translation, thus there is an important demand from the users to be able to use next generation services. From the other side, the server side, a bigger demand is approaching. The new region (e.g. Eastern Europe, Asia), where the Internet is starting to be deployed in earnest now, and the new emerging mobile/home-appliance markets, will eventually switch to IPv6, purely for transparent, global device addressability. This new wave of service providers will demand IPv6 compatible servers and services.

In this section we analyse the availability of applications that either have IPv6 support built-in or that have freely available patches to enable IPv6 support.

2.5.1.1 The IPv6 application database

There are some useful Internet sites listing general IPv6 applications, e.g.

- www.ipv6.org
- www.hs247.com
- www.bieringer.de/linux/IPv6/index.html
- www.ipv6forum.com

There are also some distribution specific sites, e.g. www.freebsd.org/ports. The Polish Linux Distribution is also quite good.

Since the information is quite scattered the TIPSTER6 project decided to implement an IPv6 ready application and patch database. The desired features of the this database were primarily that it would be simple to use, easy to search under, and also to include for the Hungarian language since the TIPSTER6 project is responsible for deployment of IPv6 in Hungary.

2.5.1.2 Implementation of the IPv6 application database

The database was implemented with freely available tools, with a Web interface designed for ubiquitous access. Perl 5 with the DBI interface was used to connect to a MySQL backend. The multilingual support was achieved with the `gettext()` library functions.

The following information is stored in the database for each application:

- User who entered the information, for delegating of privileges for further modification and later maintenance
- Name of the program
- Version of the program
- Categories the particular program falls under
- Supported operating system
- Description of the program
- Accessibility, URL where you can download the program
- E-mail address of the maintainer
- If IPv6 port exists or whether its an integrated part of the program.
- E-mail address of IPv6 port maintainer
- Accessibility, URL of the IPv6 patches
- Which information sources this record was filled out from
- Future planned date of release of IPv6 support, if not present now
- Supported protocols and RFCs of the program
- Alternatives for this program
- Comments from the entry maintainer, about usage or availability.

At this stage the software is still very much in development, but the underlying database is stable and extensible.

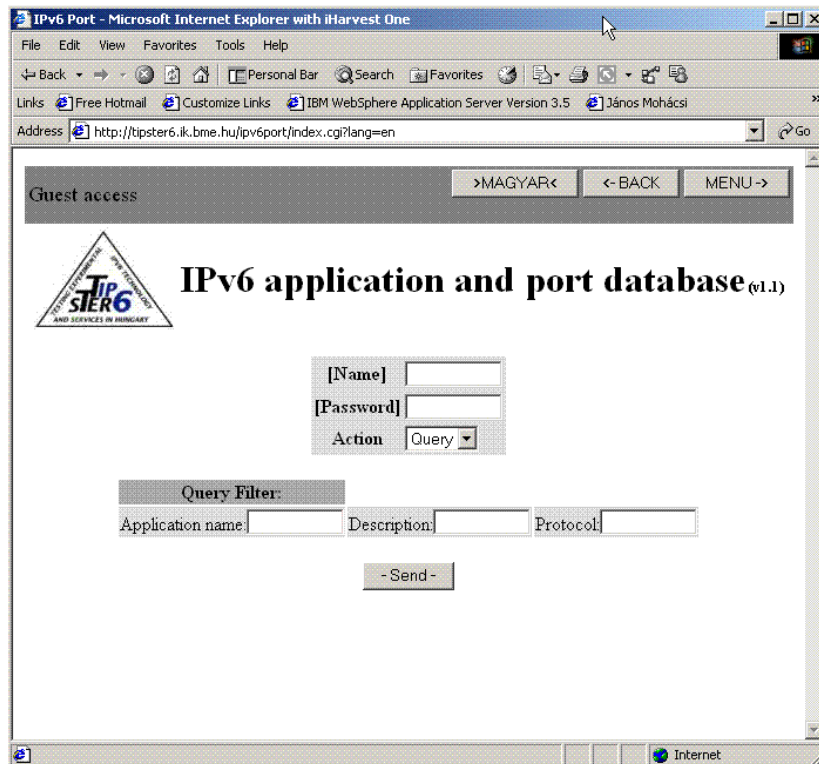


Figure 2.12: IPv6 Applications main interface

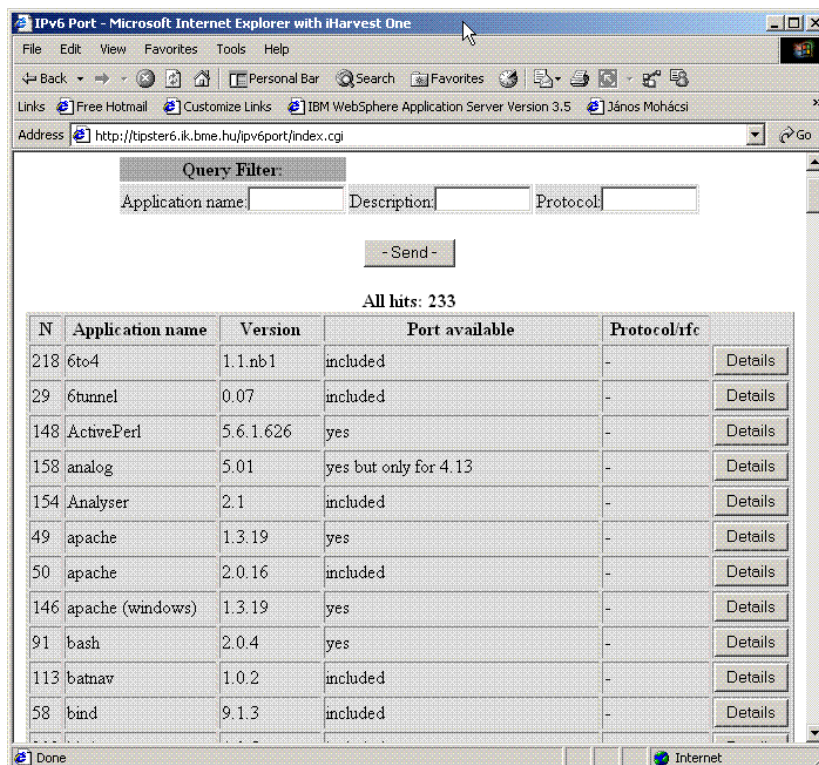


Figure 2.13: IPv6 Applications query result page

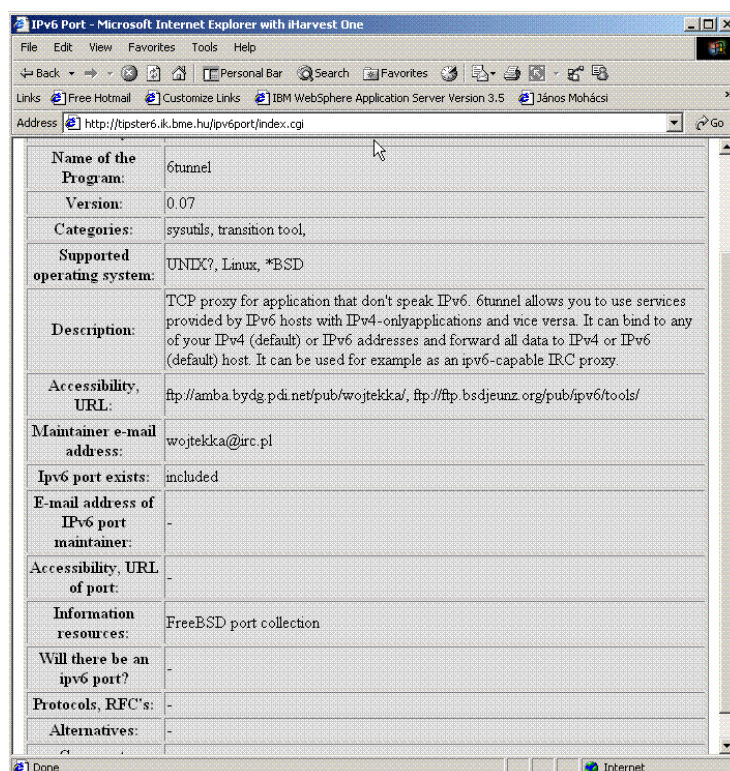


Figure 2.14: IPv6 Applications: detailed result page

2.5.1.3 Survey Results

After filling out the database, the IPv6 ready and IPv6 ported applications records were analysed. As of 13th June 2001, there were 216 packages in the database. 123 packages (57% of total) have native IPv6 support. The remaining packages (93) can be further divided in two categories:

- Patches that are in sync with current version of the package: 52 (24% of total)
- Patches that are older then the current version of package: 41 (19% of total)

The second case implies that quite a large number of applications have to be reported or to be checked against the current version of the package.

The OS support was as follows:

- UNIX: 162 (75%)
- Windows: 45 (21%)
- Linux: 172 (80 %)
- *BSD: 178 (82 %)

This implies the best IPv6 supported platform is *BSD, then Linux. Less Windows packages are available in source code, hence the porting problem. Most of the Windows ports were done by Jun Ya Kato from Waseda University.

The application categories were:

- ftp: 27
- irc: 13
- news: 7
- mail: 19
- mbone: 7
- multimedia: 11

- development:14
- system utilities: 34
- games:12
- www: 34
- testing: 22
- X11: 5
- remote login: 8
- editor: 9
- miscellaneous: 29
- DNS: 6
- transition: 8
- routing:6

We found that the most popular IPv6 application area is the WWW, then FTP, e-mail applications, and IRC. Interestingly, the JOIN traffic reports [JOINTRAF] suggest that IRC is the most popular application on their IPv6 network.

It would help porting activity, and development of new applications, if there were good documentation on methods of writing protocol independent programs. There is no automatic or semi-automatic translator for C source code. The Socket Scrubber developed by SUN [SUNIPV6] is usable but it is not comprehensive enough.

2.5.1.4 Conclusion

The IPv6 Applications database is at: <http://tipster6.ik.bme.hu/ipv6port/index.cgi?lang=en>.

The database is clearly a useful resource, but the maintenance task may be significant; there is thus a danger that like other IPv6 applications listings, it will become dated as time passes.

Discussions with the TF-NGN community led to suggestions for further work:

- Report forum, for each application that can be filled in by users, about their experiences with the package.
- Weekly mailing list about the changes.
- More search options (e.g. by category).

If more people can be allowed to add applications, the database may stay more up-to-date. It would be desirable to keep the database running and maintained as a useful resource for GTPv6 participants and other IPv6 users in general.

2.5.2 Application and porting work

A number of the NRENs offer IPv6-enabled Internet services. For example, UNINETT has production IPv6 services that include ldap.uninett.no (OpenLDAP on Linux) and www.uninett.no (IPv6 web cache running via Squid). The full list of IPv6 services offered by UNINETT is described at <http://www.uninett.no/testnett/services.html>. UNINETT is doing a number of interesting porting projects, e.g. PHP for IPv6.

JOIN offers an ftp server that is reachable with IPv6 and IPv4 ([ftp.join.uni-muenster.de](ftp://ftp.join.uni-muenster.de): IPv4/IPv6, [ftp.ipv6.uni-muenster.de](ftp://ftp.ipv6.uni-muenster.de): IPv6 only). Software archive mirrors are held on this ftp server, most being only IPv6 related material but there is also some relevant packages like linux kernels and Unix distributions (SuSe, RedHat, Mandrake). The server was recently expanded with more hard drive capacity and JOIN plans to broaden the packages offered there. The server is reachable via ftp, sftp and rsync over IPv6.

2.5.3 Videoconferencing over IPv6

The UCL conferencing “MICE” tools (vic and rat) have been ported to IPv6 [MICE]. These allow unicast or multicast IPv6 conferencing for video (vic) and audio (rat).

UoS has tested these applications successfully. However, for conferencing between sites for multiple parties, the multicast scope must be extended to those sites (for two partners, unicast remote conferencing can be used).






Figure 2.15: The vic and rat tools running unicast IPv6 (UoS – BT UK)

There is also some progress planned within a Fifth Framework project (LONG) to port the ISABEL conferencing system to IPv6.

3 SECONDARY WORK ITEMS

The secondary work items are those that the TF-NGN group felt were of lesser importance, and that therefore would receive less study in the initial reporting period leading up to this deliverable. However, since assigning the work item priorities, and leaders for each item, it is clear that, for example, there is a good deal of activity on some of the secondary items, in particular network monitoring.

It is expected that the secondary work items will be reported in more detail in an October addendum to this deliverable.

3.1 Network Monitoring

There are many properties of a network that can be monitored for utilisation, performance and effectiveness. As a group, the monitoring aspects included within GTPv6's scope include MRTG plots of IPv6 link utilisation, visual inspection of BGP4+ autonomous system (AS) routing paths, BGP "looking glass" servers, and reachability statistics (network and host uptimes).

While IPv6 patches exist for performance tools such as netperf. (<http://www.netperf.org/>), these have not been run as yet on the GTPv6 network. There is interest in the group in porting the RIPE Test Traffic Measurement [RIPE TT] servers to handle IPv6 traffic; these devices require GPS receivers for location and accurate timing information. This may also imply a need for an IPv6-enabled version of the Network Time Protocol (NTP), which is not yet currently deployed (although an RFC exists [SNTP] for Simple Network Time Protocol version 4 that includes IPv6).

3.1.1 Partner activities

JOIN has a traffic monitoring tool; this has shown that IRC is the most popular IPv6 application running on the JOIN (DFN) IPv6 network:
<http://www.join.uni-muenster.de/cgi-bin/join/trafficreport.pl>

The JOIN statistics are produced on a daily basis, and date back to July 2000. As JOIN has a very large and crowded backbone access point, the traffic might be considered as quite representative of general 6bone traffic. However, the monitoring method needs improvement, as some traffic (including http, FTP) is labelled as "generic TCP" rather than correctly by protocol.

JOIN tests site reachability (via a ping6 tool):
<http://www.join.uni-muenster.de/6bone/ping-list-e.html>

It also watches IPv6 traffic:
<http://www.join.uni-muenster.de/lab/betriebsstatistik.html>

And example of the traffic reported is given below:




Figure 3.1: JOIN traffic through the main 6bone (Telebit) router

The traffic volumes are relatively light, but it should be noted that the charts show daily traffic averages rather than peak bandwidth used at any one time. The trend is for the general volume of traffic to be growing, particularly in the last two months.

Within SWITCH, there is a BGP looking glass for the main IPv6 router:
<http://www.switch.ch/cgi-bin/lan/ipv6/look.html>

Within BME, the TIPSTER6 project has a number of monitoring tools in place. There are link utilisation statistics and peering status information:
<http://tipster6.ik.bme.hu/statistics.html>

TIPSTER6 also has its own TROUT6 tool:
<http://tipster6.ik.bme.hu/trout6/>

An example of TROUT6 output is shown below:

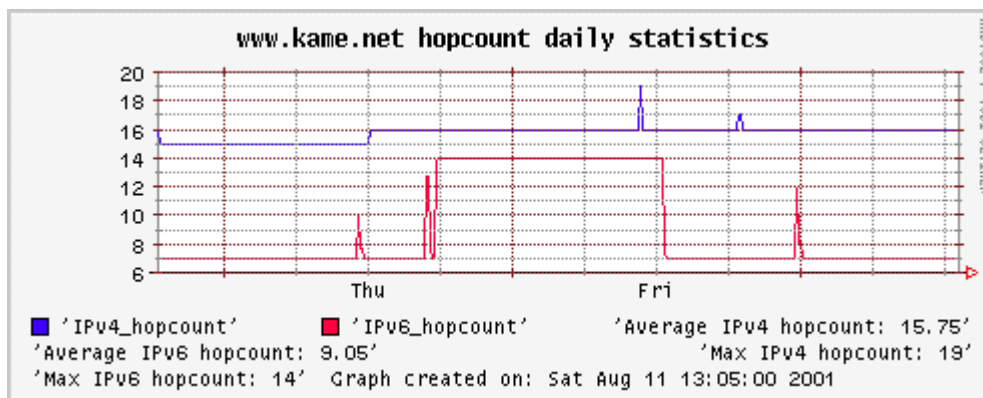


Figure 3.2: TROUT6 hop count routing information (to www.kame.net)




Figure 3.3: TIPSTER6 round trip time routing information (to www.kame.net)

The TROUT6 output is interesting because it shows both IPv4 and IPv6 hop counts and round trip times (where hosts are available over IPv4 and IPv6). The RTT values are interesting because you can compare the relative performance via each protocol, assuming the IPv4 and IPv6 versions of the target are located on the same host. If the IPv6 end-to-end link is tunnelled, the differences in RTT are more likely to be caused by poorly configured IPv6-in-IPv4 tunnels which do not best utilise the underlying IPv4 infrastructure (something that 6to4 may be better at achieving). In some cases the IPv6 links may be native, or take different routes, in which case head-to-head comparisons are not that indicative of relative performance. The hop count chart is useful for spotting changes in routing behaviour, e.g. in Figure 3.2 we can see there was a change in topology for a period of one day where the hop count from BME to www.kame.net rose from seven hops to fourteen (although the RTT in that period only rose by some 50ms).

TIPSTER6 routing information is also available based on the dynamic routing configuration observed within its IPv6 routers.

<http://tipster6.ik.bme.hu/aspath/bgp.html>

The site provides an up-to-date graphic display of the BGP4+ routing entries at BME-FSZ, 24 hours stability and odd routes circulating in the 6bone reports, generated using AS path-tree [ASPATH] as developed by CSELT in Italy. This automatically generates a set of html pages providing a graphical view of the routing paths towards the other 6Bone sites participating to the BGP4+ cloud.

RESTENA is also using AS path-tree to report BGP routing information:

<http://www.ipv6.restena.lu/bgp/>

RESTENA also has MRTG plots of link bandwidth use:

<http://www.ipv6.restena.lu/mrtg/ipv6gate/>

3.2 Multicast IPv6

The GTPv6 multicast trials have to date only been conducted within partner sites. Both UNINETT and UoS are running multicast capable IPv6 networks using the FreeBSD Protocol Independent Multicast in Sparse Mode (PIM-SM).

At UoS, the PIM-SM architecture has been used successfully for videoconferencing (using vic and rat), for an in-house MP3 multicast jukebox server (controlled via a web interface), and for a sample “space shoot-em-up” game (written for Linux with OpenGL, with game updates multicast to all players). The multicast network is implemented using FreeBSD 4.3 routers (with KAME snap), in 1U rack-mount format, with DEC quad Ethernet cards offering up to 12 IPv6 subnets, each capable of routing/forwarding IPv6 multicast packets.




Figure 3.4: Part of the multicast IPv6 PIM-SM network at UoS

The next stage of the multicast trials will involve trans-GTPv6 multicast. Initially we expect to connect UoS and UNINETT via a multicast IPv6 tunnel. Full end-to-end multicast requires multicast support on all routers between the end sites.

3.3 Wireless Access

Wireless LAN technology is advancing rapidly. With 11Mbit/s 802.11b access points and PCMCIA interface cards having replaced the previous generation of 2Mbit/s cards, the arrival of 54Mbit/s 802.11a systems (using frequency ranges that will not clash with Bluetooth) is imminent. With better performance, these systems may become more popular than fixed Ethernet, especially when users are roaming with laptop and handheld devices (such as the Compaq iPAQ, which supports 802.11b cards, albeit with a somewhat limited battery life).

UoS is deploying Lucent Orinoco 802.11b access points throughout its main Department building (8 x AP1000 access points). These run IPv4 and IPv6 (the device has to be managed via IPv4) with IPv6-only connections from at least two such points. Each AP will carry a unique routed IPv6 64-bit subnet. The IPv6 routing is currently done by servers with quad Fast Ethernet cards running FreeBSD4.3.

UNINETT has experience of wireless LANs and, interestingly, wireless MANs at the Tromso site. At Tromso, directional wireless links are used to connect IPv6-only networks within the dispersed University site. The TF-NGN meeting at Tromso in June 2001 had Internet connectivity for attendees via such an IPv6 wireless LAN.

The GTPv6 work is currently validating IPv6 access over wireless LANs, testing the different access devices, studying multicast requirements over wireless, and doing some early trials of publicly available Mobile IPv6 implementations; some results will be available come October when the GTPv6 participants may produce an addendum to this deliverable.

3.4 Multihoming

IPv6 Multihoming was one of the four study areas for TF-TANT. The studies there concluded that while a number of IETF I-D's had been published suggesting multihoming mechanisms, none were mature or ready for deployment in earnest. There is also the question of whether multihoming is something the NRENs use as a matter of course in the IPv4 world.

The situation is not that different now. The new IETF multi6 working group [MULTI6] has been formed to tackle the subject. This WG has two Internet Drafts in progress, one specifying how multihoming is done in IPv4, the other suggesting possible methods for IPv6 multihoming.

There is some difference of opinion on the best way forward. The classic approach for IPv6 is to keep the Default Free Zone (DFZ) routing table size minimal on the backbone of the Internet through heavy prefix aggregation. This implies that multihomed sites take address space from multiple providers and all hosts at such sites become multihomed. Such hosts have multiple IPv6 addresses, and thus have to make decisions on source and destination address selection, as well as handling endpoint address changes transparently while TCP sessions remain live.

In contrast, in the IPv4 world multihoming is done by advertising small address prefixes (possibly /24's) through secondary providers, causing considerable growth in the size of the DFZ routing table. There are estimates that 75% of the DFZ is represented by such /24 IPv4 networks. Were this approach taken with IPv6, the potential for DFZ growth is considerably bigger; instead of there being 2^{24} possible /24's (ignoring multicast and other prefixes) there are 2^{32} sites (of /48 size) under the 2001::/16 prefix alone, a jump of at least two orders of magnitude.

For host multihoming, with TCP transparency, some proposals have been made, e.g. "IPv6 addressing and Stream Control Transmission Protocol", an IETF Internet Draft put forward in June 2001 [SCTP], and a Masters Thesis on IPv6 Multihoming and Session Continuity by Troels Walsted Hansen [HANSEN]. There is also a new IETF Working Group emerging, called HIP, that intends to investigate the use of per-session host-based identifiers (e.g. the bottom 8 bytes of the IPv6 address) for mobility and multi-homing mechanisms (this is a variant of the old 8+8 proposal). By identifying hosts from the host part of the address, it may be possible to modify the network part of the address on the fly.

Another important issue related to multihoming is that of non-provider-based addressing, or portable address space. Part of the reason for the growth in DFZ size is the breaking of aggregation caused when large sites change provider and take their previous provider address space with them. Private addressing and NAT are seen as one method to give provider portability, but this ignores the problems imposed by NAT (lack of transparency, scalability issues, etc). In the IPv6 world it is expected that sites will renumber when moving provider, but at present, even with IPv6 autoconfiguration, the renumbering cost is high because, for example, there will still be host applications with IPv6 addresses hard-coded in them. Clearly developers and systems administrators need to consider methods to avoid such configurations.

There is scope of GTPv6 partners to test multihoming methods. Most NRENs now have both RIPE production IPv6 address space and older 6bone address space. Regarding source address selection [DRAVES], UNINETT has written a primitive patch for Linux kernels to use source addresses with longest common match to destination address. This in use at UNINETT, and also by others including CERN. A number of implementations, including KAME, are now moving to support the [DRAVES] specification, as it is near hardening to RFC status.

3.5 IPsec

One of the proclaimed advantages of IPv6 is that it mandates IPsec, yet implementations are still quite scarce. IPv6 implements IPsec by means of AH and ESP extension headers.

The two implementations under test within GTPv6 are the KAME FreeBSD (using the Racoon daemon) and the IPv6 port of Linux FreeS/WAN (with Pluto daemon). The latter was done by IABG as part of their 6INIT project work.

RESTENA has set up the two IPsec gateways, and is now beginning to test their IPv6 functionality. The installation of Racoon and FreeS/WAN took a little more time than expected, but was (relatively) straightforward.

The IPsec test results should appear in the October addendum to this deliverable.

3.6 Firewalls

One of the barriers to IPv6 deployment is the lack of commercial firewall products that can be used to “protect” a site from unwanted attention from “hackers”. While the huge EUI-64 host address space (2^{64} possible IPv6 host addresses within a single /64 network) makes port scanning in IPv6 a much smaller worry than with IPv4 (unless administrators insist on using easy-to-remember manual IPv6 addresses), there is still a requirement to protect services on hosts (e.g. SMTP or HTTP ports) that do not communicate externally from a site.

That said, in IPv6 we expect to see far greater peer-to-peer communication (due to always-on modes and the greater global addressability of devices), and we may also expect more end-to-end IPsec usage. Such factors may conflict with firewall deployment, as of course may the problem of running firewalls at Gigabit speeds. On the other hand, one might argue that it is easier to filter IPv6 packets with options (next headers) than IPv4 options.

Within the 6INIT project, UoS ran FreeBSD with its built-in ip6fw package as the firewall, though the ip6fw software was only really acting as a packet filter. With ip6fw you can specify rules to allow or deny certain packet types from entering or leaving your network on a per-interface basis. This, for example, was configured to only allow POP and SMTP traffic to certain hosts inside the network.




Figure 3.5: Using FreeBSD ip6fw for packet filtering

The various BSD implementations have mixed firewall functionality. Each variant – FreeBSD, NetBSD, OpenBSD and BSDi – has different packages supported, some with IPv6 support, some without. Even the presence of ip6fw and ipfilter (which we believe has support for stateful filtering) is not consistent across the different BSD platforms.

The TIPSTER6 project has used Linux netfilter and iptables. The Linux 2.4 kernel iptables has IPv6 extensions, and Linux 2.4 has netfilter6. We expect to produce results on the use of these tools in an addendum to this deliverable.

There is also support in Cisco IOS for IPv6 access lists, but these do not allow for stateful IP filtering (e.g. for FTP sessions).

3.7 Other items

There were some IPv6 topics not covered by the GTPv6 studies. These included IPv6 and QoS, and also DHCPv6. In the former case, it is felt that the difficult QoS issues are (relatively) IP version independent, e.g. inter-domain bandwidth brokering between DiffServ domains for end-to-end QoS. It may be that IPv6 can help, e.g. by streamlining packet classification, or innovative uses of the Flow Label field or new extension headers, but that the “political” and inter-domain issues are currently not solved for IPv4.

For DHCPv6, we await implementations. We understand that the IPv6 Forum is currently producing an RFQ for an open source, public domain DHCPv6 implementation. DHCPv6 will be important not least because it plays a vital part in a number of transition methods (e.g. DSTM).

4 INTEROPERABILITY (TAHI) TESTS

The GTPv6 scope did not originally include (formal) end system standards conformance or interoperability tests. However, the TIPSTER6 project did some good work in this area, and we report on that activity in this section.

When introducing new technology, it is important to find out the products' compliance to the standards, and their interoperability. This section describes the current state of the ongoing interoperability and compatibility tests of IPv6 implementations in the TIPSTER6 project. The results are an extension of the previous report that appeared in the TERENA Networking conference in 1998. [TNC98]

The TIPSTER6 project includes participation by the Centre of Information Technology of Budapest University of Technology and Economics Budapest, Department of Control Engineering and Information Technology of Budapest University of Technology and Economics, KFKI Research Institute for Particle and Nuclear Physics, and MATÁV PKI. Its goal is to facilitate IPv6 deployment in Hungary especially in the HUNGARNET.

Recently several IPv6 implementations have been made public. The TIPSTER6 project is concentrating on checking the implementations' adherence to standards, their basic IPv6 functions and their interoperability. TIPSTER6 selected two tools for testing IPv6 conformance: the TAHI conformance test suite [TAHI], and the Sun Packet Shell [SUNPKT]. The selection is based on availability and coverage. Both test suites have to be extended to be able to test some features of the IPv6 implementations.

In this report we are concentrating on the TAHI test results. We used the results of the extended Sun Packet Shell for checking purposes.

4.1 The Tested Systems

The majority of the tested systems are already released products, so when defects were found the vendor or the support forum of the particular product was contacted to try to resolve the problem or to provide feedback for product improvement.

- Pentium III with FreeBSD 4.2 and FreeBSD 4.3
- Sparc Ultra 5 with Solaris 8
- Pentium III with Linux 2.2.19 and Linux 2.4.3
- RS6000-170 with AIX 4.3.3

The summary of the results was presented at the TF-NGN working group in Tromsø, Norway, in June 2001.

4.2 Overview of implementations and their test results

4.2.1 *FreeBSD 4.2 and FreeBSD 4.3*

All of the BSD implementations are based on the KAME implementation [KAME], which is considered to be one of the most feature-rich and stable IPv6 implementations.

To be able to test FreeBSD under TAHI it was necessary to configure IPv6 from `/stand/sysinstall`. According to the TAHI documentation it was recommended to remove `/etc/resolv.conf` to disable DNS address resolution and to remove the IPv4 default route from the `/etc/rc.conf`.

In the case of FreeBSD 4.2 the IPv6, the ICMPv6 and the robustness tests ran correctly with the exception of one test: Routing header handling at the end node. Examining the results with tcpdump we found that the tested machine reacted correctly to the testing sequence in that particular test case, but it sent output to a different interface. We found that the KAME stack is susceptible to misconfiguration if more than one network interface exists in the tested node and it is used in stateless autoconfiguration mode. This problem cannot be considered to be a bug, since the network node cannot decide itself which interface is the default configuration interface if more than one exists. That is why an option is available for KAME implementations to define the default network interface. This problem in the FreeBSD 4.3 test (with correct setup) disappeared.

The misconfiguration caused problems with Router Solicitation, Redirect and Router Advertisement messages. In FreeBSD 4.3 some of these errors have been resolved, e.g. correction of the router/neighbour-discovery cache handling when incorrect Neighbour Advertisement packets arrive.

The PATH MTU discovery also has a problem in FreeBSD 4.2 that is mostly solved in FreeBSD 4.3.

Unfortunately Stateless Address Autoconfiguration is only working at the basic level. The more complicated autoconfiguration cases failed. The most important problem existed in Duplicate Address Detection, and it persisted in FreeBSD 4.3. This might need some more investigation, since the test results published by TAHI are virtually errorless for Address Autoconfiguration.

Two tunnel tests failed and two tunnel tests succeeded, that is reasonable since the KAME IPv6 stack only implements configured tunnels, not automatic tunnels.

	Total	OK	WARN	FAIL	Overall
IPv6 spec	37	36/37	0/0	1/0	~/✓
ICMPv6	16	16/16	0/0	0/0	✓/✓
ND	58	31/53	1/2	26/3	×/~
Stateless AC	56/57	28/31	7/9	21/17	×/×
PMTU	4	0/3	0/1	4/0	×/✓
Tunnel	4	0/2	0/0	4/2	×/✓
Robustness	4	4/4	0/0	0/0	✓/✓
Summary					~/✓

Table 4.1: Summary of FreeBSD TAHI results

Conclusion: FreeBSD for IPv6 is relatively mature and conforms well to the standards. There is a core stable release of the OS. However, if leading edge development features like NAT-PT or Mobile-IPv6 are required, then the KAME development stack (the weekly-updated “snap” kit) is high quality, but sometimes has bugs due to its ongoing development. Note that Mobile IPv6 is currently unbundled from the KAME kit.

4.2.2 AIX 4.3.3

The AIX implementation is based on the INRIA IPv6 implementation [INRIA], which was one of the most mature and stable IPv6 implementations in 1998-1999.

To able to test AIX under TAHI it was necessary to extend the TAHI suite with the necessary command and test sequences, and then to configure IPv6 from `smit` (the graphical management console). The TAHI documentation recommended removal of the `/etc/resolv.conf` file to disable DNS address resolution and also removal of the IPv4 default route.

For AIX the IPv6 standard compliance is rather good but there are some problems, including the erroneous Next Header being incorrectly reported back to the sender node, the Solicited Node Multicast Address not being handled perfectly in each case, the 65536 octet length packets not being handled correctly, and deficiencies in the handling of incorrectly ordered fragments.

ICMPv6 compliance is not perfect: most of the errors are related to link-local address handling. However there are problems with ICMPv6, the basic ICMPv6, like ICMPv6 echo, destination unreachable, and ICMPv6 error reporting not working correctly.

The IPv6 robustness tests ran without any error.

During testing Neighbour Discovery, we found a lot of problems in AIX: with Router Solicitation messages, with Redirect and with Router Advertisements. The biggest problem is that AIX sometimes does not send Neighbour Solicitation messages, it simply fills out the neighbour cache from the incoming packets. Handling the neighbour discovery messages so roughly is unexpected, since in our previous test the INRIA stack was quite good at handling them.

PATH MTU discovery also has problems in AIX 4.3.3. Only the initialisation is working. Unfortunately Stateless Address Autoconfiguration is also only working at the basic level. The more complicated autoconfiguration cases failed.

Three tunnel tests failed and one tunnel test succeeded. This is a bug of the TIPSTER6 TAHI extension though since both configured tunnels and automatic tunnels can be seen to be working correctly when tested outside of TAHI.

	Total	OK	WARN	FAIL	Overall
IPv6 spec	37	30	0	7	✓
ICMPv6	16	8	0	8	~
ND	58	10	0	48	✗
Stateless AC	57	2	18	37	✗
PMTU	4	1	0	3	✗
Tunnel	4	1	0	3	✗
Robustness	4	4	0	0	✓
Summary					~

Table 4.2: Summary of AIX TAHI results

Conclusion: AIX's IPv6 services are quite stable, but they do not confirm to the standards as well as FreeBSD, according to these TAHI test results.

4.2.3 Solaris 8

The Solaris test was run on two different kernel versions: the kernel patch version 108528-3 and the kernel patch version 108528-6, but only the latter is summarised in Table 4.3.

To able to test Solaris under TAHI it was necessary to select IPv6 as an installation option. The TAHI documentation recommended removal of the `/etc/resolv.conf` file to

disable DNS address resolution and removal of the IPv4 default route from `/etc/defaultrouter`.

For Solaris 8 the IPv6, ICMPv6 and robustness tests ran correctly. Testing Neighbour Discovery on Solaris 8 showed us that there are some problems: Neighbour solicitation messages are not sent several times, receiving Neighbour Solicitation and Router Solicitation is sometimes problematic, especially in more complicated cases. The problems that existed in handling Router Advertisement messages were almost completely corrected in the kernel patch version 108528-06. The newer version of the kernel also improves the Redirect message handling.

Unfortunately, PATH MTU discovery, Stateless Address Autoconfiguration and tunnelling worked only at the basic level.

The results are a bit unexpected, since TIPSTER6 has been using Solaris 8 with IPv6 for a long time with no serious problems. It may be a good idea to retest Solaris in a more controlled environment, to be able to distinguish the real problems from TAHI measurement errors.

We wanted to test IPsec on Solaris, but unfortunately Solaris 8 IPsec does not support asymmetrical ICMP IPsec policy.

	Total	OK	WARN	FAIL	Overall
IPv6 spec	37	37	0	0	✓
ICMPv6	16	16	0	0	✓
ND	58	37	3	18	~
Stateless AC	56	1	14	41	✗
PMTU	4	1	0	3	✗
Tunnel	4	3	0	1	~
Robustness	4	4	0	0	✓
Summary					~

Table 4.3: Summary of Solaris TAHI results

Conclusion: Solaris 8's IPv6 services are stable and generally conform to the standards. Systems should be patched to the latest kernel patch level (e.g. 108528-08 for the Sparc platform, or 108529-08 for the Intel platform)

4.2.4 Linux 2.2.19 and Linux 2.4.3

For testing Linux we used the Debian variant. We removed the default route from the `/etc/network/interfaces` file and switched of DNS resolver removing `/etc/resolv.conf`. We have tested both kernel versions with two different settings:

A) Default settings (noted with no extension):

```
net.IPv6.conf.all.autoconf = 1
net.IPv6.conf.all.accept_ra = 1
net.IPv6.conf.all.accept_redirects = 1
net.IPv6.conf.all.forwarding = 0
net.IPv6.conf.all.router_solicitations = 3
```

B) Recommended settings (noted with -2 extension):

```
net.IPv6.conf.all.autoconf = 0
```

```
net.IPv6.conf.all.accept_ra = 0
net.IPv6.conf.all.accept_redirects = 0
net.IPv6.conf.all.forwarding = 1
net.IPv6.conf.all.router_solicitations = 0
```

Under Linux we found a serious bug in the kernel, in the IPv6 fragment handling; this made the Linux kernel panic. We informed the developers via the mailing list, and the fix is now available in Linux Kernel version 2.4.4.

The IPv6 spec and ICMPv6 tests provided the best results with the 2.4.3 version kernel and recommended settings. Using these settings only a few errors remained related to multicasting and jumbograms.

Testing Neighbour Discovery, we found that the 2.4.3 kernel is better than 2.2.19, but it still contains a lot of bugs. PATH MTU discovery worked only at the basic level. Unfortunately Stateless Address Autoconfiguration hardly works at all. The tunnel tests succeeded, and also all the robustness tests were passed.

	Total	OK	WARN	FAIL	Overall
IPv6 spec	37	26/33	0/0	11/4	~/✓
ICMPv6	16	4/15	0/0	12/1	×/✓
ND	58	6/8	0/2	52/48	×/×
Stateless AC	54	1/1	8/12	45/41	×/×
PMTU	4/2 (the rest failed)	0/1	0/0	4/1	×/~
Tunnel	4	1/4	0/0	3/0	×/✓
Robustness	4	0/4	0/0	4/0	×/✓
Summary					×/~

Table 4.4: Summary of Linux TAHI results (2.2.19/2.4.3)

Conclusion: If you use Linux's IPv6 services, you should use at least kernel version 2.4.4 since it is more stable and more adherent to the standards. The USAGI project patch for Linux [USAGI] is promising since, according to their own TAHI test results, USAGI seems to solve many of the problems observed in these tests. We will seek to run TAHI tests on USAGI IPv6 Linux in the near future.

4.3 Summary

The IPv6 protocol is mature and standardised in RFCs. Some auxiliary standards are under development. The operating system vendors have finally started to support IPv6 in their products. We can safely claim that all the tested operating systems work well at the basic level, but there are some implementation problems. Based on our tests all four implementation can be used for providing IPv6 services, if recommendations are followed.

It should be noted that there may be problems in the TAHI tests themselves; this requires further study.

We are also planning to test some important systems like Tru64 UNIX, Window 2000/XP and HP-UX (as well as USAGI Linux). The details of the tests are available in the TIPSTER6 project web pages [TAHI_TESTS].

5 LIAISON WITH EXTERNAL IPV6 PROJECTS

The GTPv6 work has fed into a number of external activities and initiatives. While the primary objective of the GTPv6 studies is to move forward towards a native IPv6 deployment on GÉANT (and the connecting NRENS), the work should not be treated in isolation, and thus external collaboration is very important.

The IETF has three working groups that are focussed on IPv6 standards, namely IPng [IPNGWG] (for core IPv6 standards, which may be renamed IPv6 soon), ngtrans [NGTRANS] (for transition mechanisms) and multi6 [MULTI6] (for IPv6 multihoming). A number of NREN representatives are active within the IETF.

RIPE, along with ARIN and APNIC, sets IPv6 address allocation policies. RIPE of course has a fundamental interest in registry operations. Wilfried Woeber (ACOnet) participates in the RIPE meetings.

The IPv6 Forum [IPV6FORUM] is a market-oriented IPv6 promotion body. Some NRENS (e.g. UKERNA) have joined the Forum. The Forum includes a Technical Directorate, which discusses technical IPv6 deployment issues; Stig Venaas (UNINETT) and Tim Chown (UoS/UKERNA) are both members of the Directorate.

The European Commission has IPv6 concertation initiatives. The IPv6 Task Force [IPV6TF] is a gathering of commercial vendors and ISPs, with industrial and academic representatives, seeking to produce an IPv6 briefing paper in time for the EC Heads of State meeting in Spring 2002. The TF is covering topics such as 3G, IPv6 trial networks, and next generation applications. Tim Chown (UoS/UKERNA) is the TF rapporteur. Another EC initiative is 6LINK, a concertation of Fifth Framework IPv6-related projects (e.g. 6WINIT, LONG, Moby Dick); again Tim Chown (UoS/UKERNA) has represented GTPv6's interests in that forum.

Although it has not yet been finalised (the negotiations are ongoing at this time), the 6NET project, if it comes to fruition, will certainly have close ties with GTPv6. The Cisco-led 6NET project will provide funding and focus for NREN studies on IPv6, towards deployment of a native IPv6 pilot network between a significant number of those NRENS. However, GTPv6 will continue to be valuable for a variety of purposes, e.g.:

- Testing and deployment of other router vendors' equipment (e.g. Juniper, Hitachi).
- Bringing other NREN representatives into the GÉANT IPv6 community.
- Experiments outside the 6NET scope, e.g. this may include subjects such as multihoming or renumbering (depending on the final 6NET Description of Work).
- Acting as an "IPv6 think tank" (where 6NET is more "rigid" in its activities).

The GTPv6 group may also expect to interact with the provisional Euro6IX project in some way. Euro6IX has a commercial IPv6 focus, rather than an academic focus as per 6NET.

6 CONCLUSIONS AND FUTURE WORK

The most successful aspect of the GTPv6 work is that the activities and collaborations between the NRENS have led directly to the 6NET proposal as currently under negotiation with the Commission. However, as stated above, GTPv6 will continue to function and be a valuable activity as part of the GÉANT TF-NGN future network technology studies.

We can make some short conclusions on each GTPv6 work item:

- **Core network and routing.** The GTPv6 testbed network still uses the same structure as QTPv6 (under QUANTUM). This has proven valuable, with seven participants connecting to the core router via ATM PVCs and eight participants connecting via tunnels. Recent additions include RESTENA (Luxembourg) and POZNAN (Poland). We expect to expand this core testbed by October 2001, bringing in vendor equipment such as Juniper and Hitachi. In this document we have shown some representative NREN IPv6 network topologies; these also invariably tunnelled, with many NRENS likely to offer native IPv6 when the vendors have commercial IPv6 support in hardware (with dual-stack routers and dual-purpose links). The JOIN (DFN) deployment is probably the most extensive.
- **DNS.** There is some delay in implementation and deployment caused by the ongoing IETF debate over AAAA vs A6 records. GTPv6 has deployed an operational A6 root DNS tree, using ip6.arpa for inverse delegations. We have shown that AAAA and ip6.int results can be synthesised from this structure. More extensive A6 tests could to be run; we also need to secure an ip6.int reverse delegation for external traceroute node name reporting. At the present time it appears that the IETF will move A6 to Experimental status, thus focusing deployment on AAAA records.
- **Registries and addressing.** The GTPv6 testbed allocated 6bone addresses to participants under 3ffe:8030::/28; each participant gets a /34 allocation, which they can use as they see fit. Sections 2.1 and 2.3 describe some of the regional allocation policies. The policy for a /48 per site is good, but there needs to be clarifications as to what a site is. We would also like to see the slow roll-out policy lifted, such that NRENS are offered /29 SubTLAs instead of /35's. Moving the allocation boundary to align at a /24 or /28 may also be beneficial.
- **Transition tools.** The IPv4 to IPv6 integration and transition will occur at three levels; the core network, the NREN networks and the end-user sites (universities). The GTPv6 group has gained experience of the use of tunnels, 6to4, tunnel brokers, NAT-PT and application-layer proxies. From the experiences, guidelines for NRENS and end-user sites could be produced.
- **Applications.** BME has produced an extensive applications database, showing over 200 IPv6 applications or ports of software. A number of IPv6 services are being run by certain NRENS, e.g. web servers (Apache), mail servers (Sendmail for SMTP, qpopper for POP) and IRC chat servers. IPv6 videoconferencing has been run using the UCL vic and rat tools. Host-based operating systems show good conformance to the standards, with implementations for Solaris 8, Windows 2000/XP, FreeBSD, OpenBSD, NetBSD, AIX, and various Linux distribution variants.
- **Network monitoring.** The GTPv6 group has a number of IPv6 network monitoring services in place, including BGP looking glass pages, IPv6 link utilisation statistics (reported by MRTG), visual BGP AS routing path pages, homebrew traffic routing

tools (e.g. TROUT6), and reachability statistics (via ping-based scripts). One future activity would be to run IPv6 netperf, or to develop IPv6-enabled versions of the RIPE NCC Test Traffic servers.

- **Multicast IPv6.** At least two GTPv6 sites have run multicast IPv6 internally, using the FreeBSD PIM-SM code. This has not been extended across GTPv6 due to limitations on the intermediate router(s), but the multicast traffic could be tunnelled, or the routers upgraded/replaced with PIM-compliant routers. UoS has run multicast test packages such as a web-driven MP3 jukebox.
- **Wireless LAN access.** Some GTPv6 sites have run IPv6 over wireless LANs using the 802.11b compliant equipment (e.g. Orinoco/Lucent AP1000 access points with 11Mbit/s silver PCMCIA cards). The access points require IPv4 configuration, but will carry IPv6-only wireless LAN traffic. Next generation 802.11a equipment may prove the best option to avoid Bluetooth interference concerns.
- **Multihoming.** This has not been an active area of study under GTPv6, because the IETF methods/standards are still under discussion, and it is not clear how common multihoming is in an academic environment. The general concensus in GTPv6 is that the “small DFZ, strong aggregation” model is the one to follow, but as yet implementations of IPv6 stacks that can support transparent continuation of services across TCP sessions where links fail are rare.
- **IPsec.** GTPv6 is testing both the FreeBSD (racoon) and FreeS/WAN (Pluto) implementations. Very few IPv6 stacks support full use of the “mandated” IPv6 AH and ESP IPsec extension headers.
- **Firewalls.** There are no commercial products, but GTPv6 has operational experience to a relatively small scale of the use of FreeBSD (ip6fw) and Linux-based (netfilter, iptables) filters. Cisco IOS also supports basic IPv6 packet filtering.
- **Host-based interoperability tests.** BME ran the TAHI conformance and interoperability tests on four operating systems. In general the results were positive, and there are clearly IPv6 host platforms that can be used for day-to-day services, but the results show that all implementations have requirements for further development work. Other platforms (e.g. USAGI Linux, Tru64, Windows 2000) should be tested.

Ongoing GTPv6 work is largely a continuation of studies into existing work items. However, there are some important areas that the group wishes to develop.

International IPv6 links need further development, as do trials of new vendor equipment. In that light, the imminent deployment of an IPv6-deicated Juniper M5 router, to be hosted by RENATER in Paris on behalf of GTPv6, is welcomed. This router is planned to peer with another M5 router at the 6TAP in the USA, though the nature of the network link is not yet clear.

The group also wishes to undertake tests of the Hitachi GR2000 router series, which is now being marketed more aggressively by Hitachi in Europe. This may be a candidate router for some of the GTPv6 “backbone” functionality (and having evolved from KAME code, the router may support PIM-SM for multicast).

Although not necessarily within the scope of GTPv6, we feel it is important that more work is done in studying what required in end-user sites for IPv6-only network deployment. While

much deployment will be dual-stack, it is not until you run an IPv6-only network that all the “missing pieces” are discovered.

The GTPv6 group should also formulate recommendations for NRENs to roll out IPv6 services to universities, with the ultimate aim of native IPv6 connectivity for the universities. As discussed above, when the NREN routers can run dual-stack with hardware support for IPv4 and IPv6, this task will be made much easier. There are clear advantages to encourage universities to connect via their NRENs rather than establishing “ad-hoc” 6bone connectivity.

7 REFERENCES

For RFC documents, see <http://www.ietf.org/rfc/rfcNNNN.txt>, where NNNN is the RFC document number.

[196REV] Review of RIPE-196 Document, April 2001,
<http://www.enigma.ie/articles/global-ipv6-alteration.html>

[6BONE] 6bone, <http://www.6bone.net/>

[6INIT] 6INIT Project, <http://www.6init.org/>

[A6] RFC 2874: DNS Extensions to Support IPv6 Address Aggregation and Renumbering, July 2000, <http://www.ietf.org/rfc/rfc2874.txt>

[AAAA] RFC 1886: DNS Extensions to support IP version 6, December 1995.
<http://www.ietf.org/rfc/rfc1886.txt>

[ADDRALLOC] IPv6 Production Address Assignments,
<http://www.dfn.de/service/ipv6/ipv6aggis.html>

[ALLOC] Provisional IPv6 Assignment and Address Allocation Policy Document [ALLOC]
<http://www.ripe.net/ripe/docs/ripe-196.html>

[ASPATH] AS path-tree, <http://carmen.ipv6.cselt.it/ipv6/tools/ASpath-tree/index.html>

[AToM] Cisco AToM,
http://www.cisco.com/warp/public/732/Tech/mps/mps_learnabout.shtml

[AUSTEIN] Tradeoffs in DNS support for IPv6, R Austein
<http://www.ietf.org/internet-drafts/draft-ietf-dnsex-ipv6-dns-tradeoffs-00.txt>

[BERNSTEIN] The case against A6 and DNAME, Dan Bernstein
<http://cr.yp.to/djbdns/kill6.html>

[BLANCHET] A flexible method for managing the assignment of bits of an IPv6 address block, IETF I-D, March 2001.
<http://www.ietf.org/internet-drafts/draft-ietf-ipngwg-ipaddressassign-02.txt>

[CISCO] Cisco IOS IPv6 , <http://www.cisco.com/warp/public/732/Tech/ipv6/>

[DANTE] DANTE, <http://www.dante.org.uk/>

[DRAVES] Default Address Selection for IPv6,
<http://www.ietf.org/internet-drafts/draft-ietf-ipngwg-default-addr-select-05.txt>

[GEANT] GEANT, <http://www.dante.org.uk/geant/>

[GTPv6] GEANT Test Programme for IPv6, <http://www.ipv6.ac.uk/gtpv6/>

[HANSEN] “IPv6 Multihoming and Session Continuity”, Troels Walsted Hansen, Tromsø University, <http://www.vermicelli.pasta.cs.uit.no/ipv6/students/troels/index.html>

[INRIA] INRIA IPv6 Implementation, <http://www.ipv6.org/impl/inria.html>

[IPNGWG] IETF Ipng Working Group, <http://playground.sun.com/pub/ipng/html/>

[IPv6FORUM] IPv6 Forum, <http://www.ipv6forum.com/>

[IPV6TF] The IPv6 Task Force, <http://www.ipv6-taskforce.org/>

[JOINTRAF] JOIN IPv6 Traffic Reports,
<http://www.join.uni-muenster.de/cgi-bin/join/trafficreport.pl>

[KAME] KAME Project is a joint effort of seven companies in Japan to provide a free IPv6 and IPsec (for both IPv4 and IPv6) stack for BSD variants to the world.
<http://www.kame.net>

[MICE] UCL IPv6 Conferencing Tools, <http://www-mice.cs.ucl.ac.uk/multimedia/software/>

[MULTI6] IETF Multi6 Working Group, <http://www.ietf.org/html.charters/multi6-charter.html>

[NGTRANS] IETF ngtrans Working Group, <http://www.6bone.net/ngtrans/>

[RIPE TT] RIPE NCC Test Traffic Measurements, <http://www.ripe.net/ripence/mem-services/ttm/>

[SCTP] IPv6 addressing and Stream Control Transmission Protocol,
<http://www.ietf.org/internet-drafts/draft-stewart-tsvwg-sctpipv6-00.txt>

[SNTP] RFC 2030, Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI, <http://www.ietf.org/rfc/rfc2030.txt>

[SUNIPV6] IPv6 for Solaris, <http://www.sun.com/solaris/ipv6/>

[SUNPKT] The Sun Packet Shell, <http://playground.sun.com/psh/>

[TAHI] TAHI Project is the joint effort formed by the three Japanese organizations with the objective of developing and providing the verification technology for IPv6.
<http://www.tahi.org>

[TAHI_TESTS] TAHI test results in TIPSTER6 project. http://tipster6.ik.bme.hu/tahi_tests/

[TERENA] TERENA, <http://www.terena.nl/>

[TF-NGN] Task Force: Next Generation Networks, <http://www.dante.org.uk/tf-ngn/>

[TF-TANT] Final Report of TF-TANT, <http://www.dante.net/tf-tant/final-report.pdf>

[TNC98] J.Mohácsi, Sz. Szigeti: Testing IPv6 implementations, TERENA Networking conference, 1998, Oct, Dresden. In English <http://tipster6.ik.bme.hu>

[UNI2000] UNINETT IPv6 Report 2000, Stig Venaas,
<http://www.uninett.no/arkiv/rapport/ipv6-2000.html>

[USAGI] USAGI (UniverSAI playGround for IPv6) Project works to deliver the production quality IPv6 protocol stack for the Linux system. <http://www.linux-ipv6.org>