GOVERNMENTS ENABLED WITH IPv6

GEN6

IPv6 IN ACADEMIC NETWORKS

This project has received
funding from the
European Union´s

European
Commission

INTRODUCTION

The explosive growth of Internet and the fast depletion of the conventional IPv4 addresses lead to the development of IPv6 as a new network layer protocol. While IPv6 has been around since 1998, several intermediate solutions to solve the Internet addressing problem, namely using private addresses and the Network Address Translation protocol (NAT), refrained its wide deployment. However, in January 2011, IANA distributed the last few blocks available at the top-level domain. With this event, IPv6 became a strategic and mainstream topic for governments and telecommunication companies. In the last few years, we assisted to a huge progress on IPv6 deployment at world level. While a lot of transition effort is still required, as of March 2015 Google reports about 6-7% of IPv6 based requests at world level, up from a residual 0.5% penetration in mid 2011. This represents a remarkable 14-fold increase in just 4 years.

Academia usually is, by its own nature, a pole of innovation and an early adopter of new technologies. This was also partially true regarding IPv6. The GÉANT network is the pan-European research and education network that interconnects Europe's National Research and Education Networks (NRENs). The GÉANT backbone fully supports IPv6 since 2002, as well as the majority of European NREN networks. These dates preceded by far the adoption of IPv6 by major telecommunication companies, which only in the last few years started to address the transition. However, IPv6 penetration at the Academic, institution and user level is quite variable and, quite surprisingly, does not even seems to match the increased penetration of IPv6 at the global level. Informal figures suggests that the total weight of IPv6 traffic in the GEANT backbone is around 3%, a few points below global traffic.

The goal of the GEN6 project was to increase awareness and promote the adoption of IPv6 at Government and Public Administrations Services. The Academic Pilot was developed as part of this effort and aims to promote good practices for IPv6 deployment at academic networks, as well as promoting the on-line monitoring of effective IPv6 penetration at user level in academic networks.

WHY IPv6 NOW (ESPECIALLY IN ACADEMIA)?

Deploying IPv6 today has become a reality that every organization „must" face. There is a very strong motivation: there are not as many internet addresses in the old version of the Internet Protocol (IPv4) as there are people on the planet, and we are running out of addresses. IPv6 provides trillions of addresses that will enable the internet to grow, for all practical purposes, indefinitely.

Some factor reinforce this need of IPv6 deployment in academia. First of all, it prevent increased costs, as well as disruption of the web sites. Moreover, the number of „connected" devices (tablets, smartphones, laptops, ...) continuously grows and requires even more IP addresses.

After a 20 years development, IPv6 is ready to deploy today. Most competitors are already doing IPv6 (major Internet Service Providers, GÉANTs research network, web companies, etc...).

CONSIDER THE CHANGES IN YOUR CURRENT IPv4-INFRASTRUCTURE

There is no need to restructure the complete network for IPv6, if:

- the security areas are well-designed (edge, DMZ, backbone...);
- the IPv4 network is already split into subnets;
- the design is scalable.

Otherwise, restructuring should be carried out before starting with IPv6.

Here are some guidelines that can help you to redesign the network infrastructure. Maybe you have to split a „one big network" into subnets, in order to isolate broadcast domains and spanning-trees. The subnets must reflect an internal structure, and must be clearly defined. Then you can use V-LANs to manage the subnets. Natting should be avoided (natting with IPv6 is not recommended). In the routing area, dynamic routing protocols should be deployed in order to build redundancy.

DEFINE THE IPv6 STRATEGY

The IPv6 strategy involves the target architecture, a high-level address plan and a high-level security concept (which must be aligned with each other).

It is easier to roll out IPv6 over a structured IPv4 network to avoid huge amounts of downtime during implementation and general confusion by trying to change too many things at the same time.

The transition to IPv6 gives an organisation the opportunity to fix problems in its current environment. The steps are:

1. evaluate network equipment and applications that have to support the transition to IPv6 (hardware, services, management, applications, workstations, etc...) and identify the items that are not IPv6-compliant

2. figure out the IPv6 subnets needed. It is usual to plan as many IPv6 subnets as IPv4 ones existing in the network. Of course, leave a marge for future expansion (IPv6 prefixes are larger than IPv4 ones, you will have enough place).

3. get an IPv6 prefix for your organization (from your ISP, or under certain conditions directly from RIPE).

4. create an address allocation plan. The concept must be well structured, simple, scalable and clear. The easiest, but least flexible solution is to make block address assignment, starting from the beginning of the organisation's allocated IPv6 block. For future needs, it is possible to keep some prefixes free. The prefix or network part of an IPv6 address should always be 64 bits (/64).

5. agree on the way you will manage addresses (DHCPv6/DHCPv6, SLAAC, manual).

6. design the routing schema infrastructure. Above all, the routing protocols in use (static, OSPF, BGP, ...) and the adaptations needed to enable IPv6 routing are to be considered.

IPv6 NETWORK TRANSITION MECHANISMS AND IPv6 DEPLOYMENT STRATEGIES

Three transition possibilities are given for the configuration of IPv6 in a network: dual stack, tunnelling and translation.

The decision of the possibility to adopt can only be done regarding the global infrastructure and the future concept that is decided on. Dual stack has the advantage that IPv4 can live parallel to IPv6 as long as it is needed.

Then, a strategy must be adopted for the implementation of IPv6. One example is to deploy IPv6 in the core first, then on the clients. On the other hand, the deployment can be based on geographical implementation.

The strategy to adopt depends on the infrastructure and layout of the organization.

SOME HINTS ON NETWORK HARDWARE AND SERVICES

After taking the inventory, all the products should be checked and evaluated for IPv6 compliance. Obtain a list of implemented RFCs from the vendor, check it against your requirements, and then test in your lab whether the features that are critical to your plans work as expected.

Especially for layer-3 equipment, several protocols and features must be checked for compliance (e.g. TCP, UDP, ICMP protocols, access-lists, policy-based routing, dynamic routing protocols, etc...).

Most network services like DNS, DHCP, NTP, VPN (OpenVPN, IPSec)... are today IPv6 ready and can be enabled after a few configuration work will be done.

IPv6 MULTICAST

Multicast is an efficient way to deliver data whenever one source is distributing the same data to multiple receivers. Multicast in IPv6 is not much different from IPv4. There are only new concepts you need to be familiar with:

- Multicast Listener Discovery (MLD): MLD is the equivalent to IGMPv2, defined for IPv4. MLD messages are carried in ICMPv6 packets.
- MLD snooping: is similar to IGMP snooping proposed for IPv4.

NEIGHBOUR DISCOVERY: THE BASIS FOR ALL IPv6 IMPLEMENTATIONS

Neighbour Discovery is a protocol that allows different nodes on the same link to advertise their existence to their neighbours and to learn about the existence of their neighbours.A router periodically sends out router advertisements from each of its multicast interfaces, announcing its availability. Hosts listen to these advertisements for address auto-configuration and discovery of link-local addresses of the neighbouring routers.

Neighbour Discovery for IPv6 replaces the following IPv4 protocols: router discovery (RDISC), Address Resolution Protocol (ARP), and ICMPv4 redirect.

SOME CONSIDERATIONS ABOUT ROUTING PROTOCOLS

- Using static routes in IPv6 is similar to configuring static routes for IPv4.
- IPv6 RIP or RIPng (RIP Next Generation) generally is very similar to RIP (for IPv4) and functions in the same way, offering the same benefits.
- EIGRP for IPv6 works in the same way as EIGRP IPv4 where they can be configured and managed separately.

- Many of the OSPF for IPv6 (OSPFv3) features are the same as in OSPF for IPv4 (OSPFv2). The major differences between both protocols are:
  - OSPFv2 and OSPFv3 are completely separate instances
  - OSPFv3 uses multicast for communication.
  - OSPFv3 uses link-local addresses for routing.

- IS-IS (Intermediate System to Intermediate System) in IPv6 works in the same way and offers many of the same benefits as IS-IS in IPv4.

- Multiprotocol BGP is the supported Exterior Gateway Protocol (EGP) for IPv6. Multiprotocol BGP extensions for IPv6 support the same features and functionality as IPv4 BGP.

IPv4/IPv6 routing protocol classification

	Interior gateway protocols				**Exterior gateway protocols**
	Distance vector		**Link state**		**Path vector**
IPv4	RIPv2	EIGRP (Cisco prop.)	OSPFv2	IS-IS (Cisco prop.)	BGP-4
IPv6	RIPng	EIGRP for IPv6 (Cisco prop.)	OSPFv3	IS-IS for IPv6 (Cisco prop.)	MBGP

WLAN AND IPv6

In order to enable wireless IPv6 client connectivity in the Wireless LAN, the underlying wired network must support IPv6 routing and an address assignment mechanism, such as SLAAC or DHCPv6.

Access Points need connectivity to the IPv6 router. On layer-2, there is nothing specific to do for carrying IPv6 traffic. When layer-2 filters are set on the Access Points, they must allow for IPv6 traffic.

WLAN controllers transparently forward IPv6 packets.

NETWORK MANAGEMENT OVER IPv6

Every network operation centre has its tools and protocols to manage and monitor its network. All these items must be gathered and in the long term will have to migrate to IPv6.

SNMP

Today, many network vendors (6WIND, CISCO, HITACHI, JUNIPER etc.) support SNMP over IPv6 and can be monitored in an IPv6-only environment. Equipment not supporting SNMP over IPv6 may be managed over IPv4, as most IPv6 networks are running dual-stack. The number of SNMP applications able to poll remote SNMP agents over IPv6 remains low.

Flow monitoring (Netflow)

NetFlow for IPv6 is based on NetFlow Version 9 and works by identifying packet flows for ingress IP and IPv6 packets.

SDN and OpenFlow

Even if a lot of network operation centres do not implement OpenFlow in their current IPv4 network today, OpenFlow has been IPv6-ready since version 1.3.

SECURITY ASPECTS

For the success of IPv6 deployment in networks, it is important that the IPv6 deployments are secure and are of a service quality that equals that of the existing IPv4 infrastructure.

Inbound filtering of Martian packets makes sense on the internet uplink to prevent malware from using these as source address. To prevent internal and external misuse of these prefixes, the Martian prefixes can also be null-routed (e.g. a null route is propagated to the network via a dynamic routing protocol).

Martian IPv6 packets

Address block	Present use
::/128	Node-scope unicast unspecified address
::1/128	Node-scope unicast loopback address
::ffff:0:0/96	IPv4-mapped addresses
::/96	IPv4-compatible addresses
100::/64	Remotely Triggered Black Hole addresses
2001:10::/28	Overlay Routable Cryptographic Hash IDentifiers (ORCHID)
2001:db8::/32	Documentation prefix
fc00::/7	Unique local addresses (ULA)
fe80::/10	Link-local unicast
fec0::/10	Site-local unicast (deprecated)
ff00::/8	Multicast [16] (Note: ff0e:/16 is global scope and may appear on the global internet.)

As a lot of IPv6 features are based on ICMPv6 packets (Neighbour Discovery, PMTU discovery, etc…), misused nodes can result in denial of service or man-in-the-middle attacks. The recommendations in the table below allow for the propagation of ICMPv6 messages needed to maintain functionality of the network, but drop messages posing potential security risks. These basic rules should be configured on the firewall.

Some mechanisms exist to grant security at the first hop, especially in the IPv6 context:

- Router Advertisement Guard (RA Guard)
- SEcure Neighbour Discovery (SEND)
- Anti-spoofing is configured the same way for IPv4 as for IPv6.

APPLICATIONS AND LOAD-BALANCERS

Most of the web, mail, cloud services as well as load balancers are today IPv6 ready. You have just to enable the IPv6 connectivity of the service and do a few configurations to provide IPv6 for your applications.

TESTING AND DEVELOPING AN IMPLEMENTATION PLAN

To become familiar with IPv6, it can be a good idea to set up a test bed with some key features of the local infrastructure, as for example, a border router, a firewall system, a backbone router and a switch, some clients (Windows, Linux, ...) and some servers (web server, mail server, DNS...).

Once this "little" test bed is set up, start to configure some IPv6 networks (transit networks, client network), some firewall rules, access lists and play with routing protocols. Depending on your implementation plan, you can test dual-stack or tunnelling.
With intensive tests in this set-up, you will get a feeling of what needs to be done in your local infrastructure.

The next challenge will be to develop an implementation plan. Regarding the inventory of your equipment, the design you have worked out, the transition mechanism you will adopt, you will be ready to begin with the IPv6 implementation.

IMPLEMENTING IPv6

You should already have got an IPv6 prefix (from your ISP, or under certain conditions from RIPE NCC). Afterwards, you can configure your equipment according to the strategy you have chosen (core/border, geographical infrastructure, ...). The configuration should be described in the handbook of your equipment vendor. Once basic connectivity is guaranteed, you can enable IPv6 for central services (DNS, VPN, web, mail, etc...).

TROUBLESHOOTING IPv6

Troubleshooting IPv6 should be as easy as IPv4. Some tools are different depending on the platforms.

Tools and commands for troubleshooting IPv6 on Linux clients:

- ping6
- traceroute6
- tracepath6
- mtr -6
- ip -6
- host
- tcpdump/Whireshark

Tools and commands for troubleshooting IPv6 on Windows clients:

- ipconfig /all
- netsh
- tracert -6
- pathping
- Whireshark

Troubleshooting IPv6 on Cisco routers:

- ping
- traceroute
- show ipv6 route
- show ipv6 access list
- show ipv6 interface

MONITORING IPv6

Netflow

Netflow is a flow-based traffic accounting protocol developed by Cisco Systems. The latest version, Netflow v9, is used as a basis for the IPFIX (IP Flow Information eXport) protocol that is currently being standardised in the IETF (RFCs 5101 and 5102). Only version 9 of Netflow is designed to export IPv6 flows towards the Netflow collector.

Many vendors other than Cisco provide an equivalent technology on their routers and switches, but some use a different name for the technology, probably because Netflow is thought to be a Cisco trademark. Some examples are: Jflow or cflowd for Juniper Networks, Cflowd for Alcatel-Lucent, sFlow, etc… sFlow vendors include: Alaxala, Alcatel Lucent, Allied Telesis, Arista Networks, Brocade, Cisco, Dell, D-Link, Enterasys, Extreme, Fortinet, Hewlett-Packard, Hitachi, Huawei, IBM, Juniper, LG-Ericsson, Mellanox, MRV, NEC, Netgear, Proxim Wireless, Quanta Computer, Vyatta, ZTE, and ZyXEL.

The Netflow collector is responsible for reception, storage, and pre-processing of flow data received from a flow exporter.

There are some free Netflow collectors with IPv6 support:

- flowd NetFlow collector
- -ntop nProbe™ v7
- nfdump and nfsen
- PMACCT Netflow

Commercial Netflow collectors with IPv6 support are available from:

- Cisco
- Ipswitch WhatsUp
- SolarWinds
- etc…

Argus

Argus is a system and network monitoring application, which has included IPv6 support since version 3.2. It will monitor nearly anything it is asked to (TCP and UDP applications, IP connectivity, SNMP OIDS, etc). Due to the fact that most of the testing modules are written in Perl, IPv6 functionality is included in most of them.
(http://argus.tcp4me.com)

ntop

ntopng, the next generation version of the original ntop, is a network traffic probe that shows the network usage similar to what the popular top Unix command does.

ntop now fully supports IPv6 thanks to the effort of INRIA within the WP6 framework of the 6NET project. ntop can collect and make stats available on IPv6 flows and hosts in the network. Moreover, IPv6 flows can be exported with netflow v9 from with the nProbe feature and merged with the ntop tool. ntop can also act as a netflow v9 collector for IPv6 flows exported from other equipment (e.g. a CISCO router or another ntop application).
http://www.ntop.org/

IPv6 launch day

Organised by the Internet Society, the World IPv6 Launch on 6 June 2012 was intended to motivate organisations across industry – including internet service providers (ISPs), hardware makers, and web companies – to prepare for and permanently enable Internet Protocol version 6 (IPv6) on their products and services, as Internet Protocol version 4 (IPv4) address space runs out.
http://www.worldipv6launch.org/

MONITORING IPv6 SUPPORT IN ACADEMIC NETWORKS

Monitoring the support of IPv6 at academic level provides an immediate feedback which may contribute to increase IPv6 awareness among University and ICT governance.

As a first part of the Academic Monitoring task, it was performed a first assessment of IPv6 support in the backbone and central service of main European Universities. This analysis encompassed the test of web, DNS and mail services of the central services of a subsample of significant Universities across Europe. Considering the high number of Higher Education Institutions (HEI) in Europe, the scope of this analysis was for now limited to a sample of 177 of these institutions that make part of the top 400 in the Times Higher Education World University Rankings. The result of this assessment is summarized in the table below. The row "Partial IPv6 support" means that IPv6 addresses are available but no service is available, or that only part of known DNS and mail servers of a specific academic domain have IPv6 support.

April 2015

	WWW	**DNS**	**Mail**
Full IPv6	**33** *(18.6%)*	**30** *(16.9%)*	**44** *(24.9%)*
Partial IPv6	**0** *(0.0%)*	**91** *(51.4%)*	**7** *(4.0%)*
IPv4 only	**144** *(81.4%)*	**56** *(31.6%)*	**126** *(71.2%)*

While these figures seem promising and positive at a first glance, a more detailed analysis is advisable. A quite higher figure is seen for partial support of DNS services, possibly due to a high number of secondary name servers being hosted in NREN servers or other ISPs having IPv6 support. A quick analysis of some academic mail services in this subset that do not have web or NS IPv6 support shows that, in most of these cases, the mail service is hosted in NREN servers or external commercial operators with dual-stack support, which explains the higher level of IPv6 adoption in mail services. While less usual, the support of IPv6 in web services may also be slightly inflated by academic web sites hosted externally by operators with IPv6 support.

Up to date figures of this analysis may be found on line at
https://devpub.labs.nic.cz/ipv6-smt-new/country/_uni/.

MONITORING IPv6 DEPLOYMENT IN ACADEMIC NETWORKS

While monitoring IPv6 support at central services of main Universities provides a clue to the commitment of ICT departments and corporate governance to IPv6, it may not provide an accurate estimate of IPv6 adoption at user level. In fact, IPv6 support in central services does not often reflect the level of effective IPv6 deployment in distribution and access networks.

To obtain more accurate figures regarding IPv6 penetration, it was developed a small javascript probe, adapted from a previous version by the IPv6 observatory. In order to obtain accurate results, this probe should be installed as widely as possible in websites with popular content in the Academic community. When the end user accesses the website, the browser executes the probe code, and a few HTTP GET request are sent to some controlled virtual test web sites with different types of IPv6 support. Each one of these accesses simply downloads the smaller transparent image one may create (1px by 1px). This small image is designed in order to imply the effective access to the test sites, but in fact does not introduce any perturbation to the graphic content of the web site, being unnoticeable by the end user. The results of the several HTTP requests, which are asynchronous, fast and unnoticeably by the user, are logged in raw databases for later processing. The overall architecture of this model is represented below:

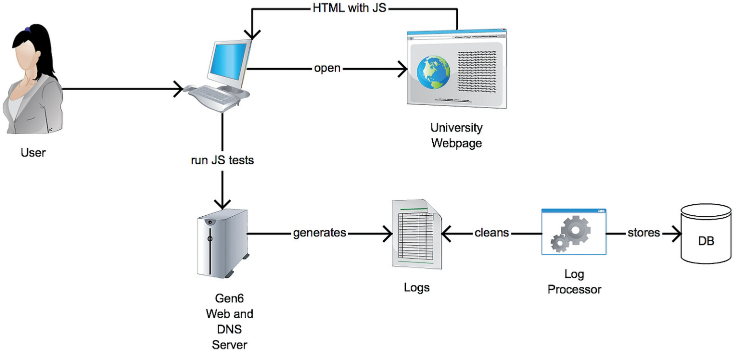

The installation of the probe was first promoted in main European University websites, since these are institutions open to research and that, by their own nature, do have a high number of accesses from the academic community (from its own population and also from other Universities). However, one must emphasize that the location of the probe itself is irrelevant for the tests performed in the scope of this activity. In fact, IPv6 support is tested between the user platform and the IST test server, irrespective of the probe location, as depicted in the figure below.

During the post-processing phase, each logged access is retrieved and the source IP analyzed trough a reverse DNS and lookup to the WHOIS database, in order to identify the source network. If it falls within a monitored Academic network, the University is identified and the associated statistics are updated. The results of this monitoring task are available online and updated daily at

https://gen6.tecnico.ulisboa.pt/

As explained before, by default this list includes only the subset of European Universities that make part of the 400 in the Times Higher Education World University Rankings. However, if your University is not included in this default subset but you want to be included in this analysis, please contact us at gen6@tecnico.ulisboa.pt. We will be glad to help and include your institution in this monitoring analysis.

DISCUSSION AND RECOMMENDATIONS

While IPv6 support in the backbone of the research and academic networks dates back to 2002, this early support at the backbone does not translates on a wide penetration of IPv6 at the end user level in academic networks. In spite of the excellent dissemination work and IPv6 promotion made both by GÉANT and national NRENs, an extra effort is still required to fully promote IPv6 on academic networks.

This document describes two main contributes that may help to achieve this goal:

(1) A comprehensive set of recommendations and best practices for deployment of IPv6 in academic networks;

(2) A monitoring system that provides, at the same time, a framework for assessment of IPv6 support in main Universities across Europe as well as a tool for auto assessment of adhering

WHERE TO FIND ADDITIONAL INFORMATION

http://www.ipv6forum.com/	World-wide consortium of Internet vendors aiming to promote IPv6. Includes mailing lists, event listings, technical information, and links.
http://www.6net.org/	A lot of useful documents about IPv6 deployment
IPv6 Essentials 3rd Edition O'Reilly	Vendor-independent IPv6 book
https://www.ietf.org/rfc.html	RFC repository
http://en.wikipedia.org	Free encyclopedia, every definition you need is there
https://devpub.labs.nic.cz/ipv6-smt-new/country/_uni/.	IPv6 support in central services of main universities
https://gen6.tecnico.ulisboa.pt/	IPv6 penetration in academic networks.

This booklet is prepared by Karlsruhe Institute of Technology TÉCNICO LISBOA

This work was in part supported by the European Commission as part of the project "Governments Enabled with IPv6" (GEN6). GEN6 is about stimulating EU-wide deployment of IPv6 by means of best practices and guidelines in existing eGovernment infrastructures.

Authors

Reinhard Strebler (KIT, Germany)
Aurelie Reymund (KIT, Germany)
Fernando Mira da Silva (IST, Portugal)
Jorge Matias (IST, Portugal)
David Duarte (IST, Portugal)
José Costa (IST, Portugal)
Luís Guerra e Silva (IST, Portugal)

This booklet is based on the deliverables in work package 3 of the GEN6 project. All participants in this working package have contributed to this booklet indirectly. They are not explicitly mentioned here. Details on the requirement analysis made and the involved partners and authors can be found on the projects web site at the category publications - deliverables.

Contact

To get in contact with the GEN6 project or the partners please contact us.

info@gen6-project.eu

www.gen6-project.eu

This booklet is part of a series of information on IPv6 transition in eGovernment. See www.gen6-project.eu/publications/booklets/ for further available booklets. Booklets already published:

- Smart communication solutions in emergency situations
- Energy efficiency in school networks with IPv6
- IPv6 application in the road domain
- Addressing and transition from IPv4 to IPv6 in government networks
- Why are governments on IPv6? Start your IPv6 project right now
- IPv6 standards and RFCs - what profiles can do
- Secure election infrastructures based on IPv6-clouds
- A National-level IPv6 addressing concept for the government
- IPv6 implementation in existing egovernment infrastructures

Layout by Citkomm, all photographs © 2015 by fotolia.com