

Zeitschrift: Comtec : Informations- und Telekommunikationstechnologie = information and telecommunication technology

Herausgeber: Swisscom

Band: 81 (2003)

Heft: [1]: A collection of publications of Swisscom Innovations from 2003

Artikel: Towards the full deployment of IPv6 : the experience at Swisscom Innovations

Autor: Schmid, Andreas / Gysi, Martin / Schädler, Michael

DOI: <https://doi.org/10.5169/seals-876728>

Nutzungsbedingungen

Die ETH-Bibliothek ist die Anbieterin der digitalisierten Zeitschriften auf E-Periodica. Sie besitzt keine Urheberrechte an den Zeitschriften und ist nicht verantwortlich für deren Inhalte. Die Rechte liegen in der Regel bei den Herausgebern beziehungsweise den externen Rechteinhabern. Das Veröffentlichen von Bildern in Print- und Online-Publikationen sowie auf Social Media-Kanälen oder Webseiten ist nur mit vorheriger Genehmigung der Rechteinhaber erlaubt. [Mehr erfahren](#)

Conditions d'utilisation

L'ETH Library est le fournisseur des revues numérisées. Elle ne détient aucun droit d'auteur sur les revues et n'est pas responsable de leur contenu. En règle générale, les droits sont détenus par les éditeurs ou les détenteurs de droits externes. La reproduction d'images dans des publications imprimées ou en ligne ainsi que sur des canaux de médias sociaux ou des sites web n'est autorisée qu'avec l'accord préalable des détenteurs des droits. [En savoir plus](#)

Terms of use

The ETH Library is the provider of the digitised journals. It does not own any copyrights to the journals and is not responsible for their content. The rights usually lie with the publishers or the external rights holders. Publishing images in print and online publications, as well as on social media channels or websites, is only permitted with the prior consent of the rights holders. [Find out more](#)

Download PDF: 03.08.2025

ETH-Bibliothek Zürich, E-Periodica, <https://www.e-periodica.ch>

Towards the full Deployment of IPv6 – the Experience at Swisscom Innovations

After many years of testing in research networks, the worldwide deployment of the Internet Protocol version 6 (IPv6) in commercial networks is now taking place. Asia (especially Japan and South Korea) is leading the way. Europe is following with a massive investment in IPv6-related R&D projects. Recently, the USA has started to seriously consider IPv6, led by the Department of Defense (DoD) that is planning to execute the full transition of the entire enterprise IP network to IPv6 by 2008.

Swisscom Innovations has been running IPv6 test networks for quite a while, and currently the operational Intranet is IPv6-enabled. This article briefly introduces an IPv6 transition scenario and describes experiences gained from the deployment of IPv6 in our Intranet.

The programme "Future Network Services" explores emerging technologies enabling wired and wireless, fix and mobile broadband services. A multitude of access technologies will coexist in the near future. Customers will be able to access the network, including voice services, through end devices supporting several access technologies, as for example DSL, GSM, GPRS, UMTS, WLAN and Bluetooth. With its Innovation Programmes, Swisscom Innovations follows the objective of recognising early on the impact of technological developments, finding new business opportunities, promoting technical synergies and developing concrete innovation proposals. Further, the expertise built up enables active engineering support of business innovation projects.

Today's Internet is based on the Internet Protocol version 4 (IPv4). Applications like the worldwide web (WWW) and email enabled the Internet to grow at huge rates. In the early 1990s, it became clear that IPv4 would not be

the end-to-end paradigm. Innovation in Internet applications was severely hindered by the introduction of NAT. It is likely that the success of the WWW would not have been possible if NAT had been introduced to the Internet during early 1990.

ANDREAS SCHMID, MARTIN GYSI AND
MICHAEL SCHÄDLER

able to sustain this rapid growth. Exhaustion of available IPv4 addresses was predicted to happen around 1995. Standardisation of a new version 6 (IPv6) of the Internet Protocol became eminent. However, mechanisms like Network Address Translation (NAT), together with the use of private IPv4 addresses and Classless Inter-Domain Routing (CIDR) extended the life of IPv4. Unfortunately, NAT breaks a fundamental design rule that made the Internet that successful:

IPv6 deployment is possible today and can be successful if planned in phases. This paper describes the approach Swisscom Innovations has taken for the IPv6 deployment in the operational Intranet and proves that commercial IPv6 deployment is possible. This is of major importance for the Internet community because IPv6 re-establishes the end-to-end paradigm and therefore lays the ground for the next innovation boost in Internet applications.

The standards of the IPv6 protocol suite were finalised during 1998. In the following 2–3 years little deployment happened with the exception of the global IPv6 test-bed "6bone". Missing customer need and lack of commercial implementations were the main factors why IPv6 was not successful at that time. This situation is currently changing (see e.g., the initiative of the USA Department of Defense [1] or the recent Comtec article by Silvia Hagen [2]). Pre-commercial IPv6 services are now offered all over the world. Most ISPs and operators are planning the introduction of IPv6 in their networks and services. Swisscom Innovations has built up extensive knowledge and practical experience in IPv6 over the last years. The time is now right to deploy IPv6 in an operational environment.

How IPv6 can coexist with IPv4

IPv6 is not inherently compatible with IPv4. Consequently, many mechanisms have been standardised to enable interworking. IPv4 and IPv6 networks will coexist for a very long time due to the huge number of IPv4 networks, hosts and applications installed to date. Different scenarios for coexistence of IPv4/v6 networks and hosts are possible where three categories of transition mechanisms are employed:

1. *Dual-stack*: hosts have both IPv4 and IPv6 stacks and are able to communicate on either protocol.
2. *Translation*: devices, called translators, are placed between IPv4 and IPv6 networks. These devices "translate" between the two versions of the Internet Protocol. With this solution, application layer gateways are often needed in addition to the translators in the network layer.
3. *Tunnelling*: tunnels are used to transport data of one protocol over a network of another protocol. In the IPv4 to IPv6 transition, IPv6 in IPv4 tunnels are widely used to interconnect IPv6 islands/hosts over IPv4 networks. Figure 1 depicts how tunnelling works.

Configured Tunnels versus Automatic Tunnels

A host performing IPv6 in IPv4 tunnelling sends packets with a specific IPv6 destination address to an associated IPv4 address (tunnel endpoint). The relationship between IPv4 and IPv6 addresses can be set up automatically or by manual con-

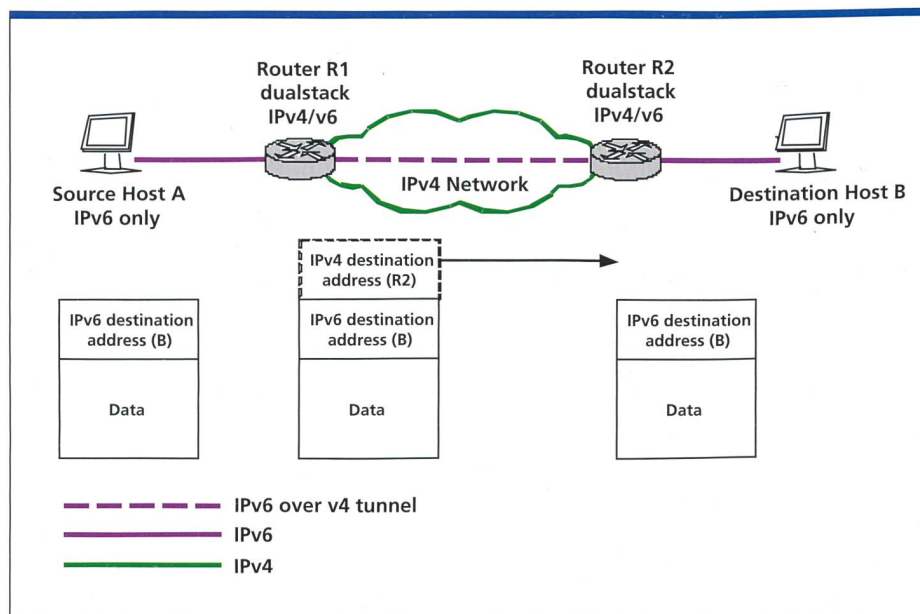


Fig. 1. Tunnelling example of a router-to-router tunnel.

figuration. Therefore, there is a distinction between configured and automatic tunnels. Using configured tunnels does not scale well due to the manual configuration needed.

For automatic tunnels, the IPv4 destination address is not known in advance and has to be derived from the IPv6 destination address of a packet. Therefore, the IPv4 destination address has to be included in the IPv6 address in one way or another. The most popular automatic tunnelling mechanisms are 6to4 [3] and ISATAP (Intra-Site Automatic Tunnel Addressing Protocol) [4]. 6to4 has been developed for use in the global IPv4 Internet where a host has a public IPv4 address and wants to tunnel over the public IPv4 Internet to other IPv6 clouds. In contrast, ISATAP is designed for use in enterprise networks where hosts with private IPv4 addresses want to reach the IPv6 Internet. ISATAP is described in detail in the following paragraph.

ISATAP

ISATAP maps the IPv4 address into the IPv6 address as shown in figure 2. To enable a network for ISATAP, an ISATAP gateway has to be installed that is connected to both the IPv6 Internet and the IPv4 private Intranet. It serves as default gateway for all IPv6 traffic from the Intranet. Dual-stack nodes use ISATAP IPv6 in IPv4 tunnels terminated by the ISATAP gateway for IPv6 connectivity. To do so they need to find the IPv4 address of the ISATAP gateway. A simple mechanism – a special entry in the DNS server for the ISATAP gateway – is used in this situation. ISATAP clients query the DNS server for the address of the ISATAP gateway and send a router solicitation to this address. The server replies with a router advertisement that contains the globally routable IPv6 prefix, allowing the enquiring clients to auto-configure their IPv6 addresses.

Microsoft is the driving force behind ISATAP and has it implemented in Windows XP SP1, Windows 2003 server, CE .NET 4.1 and Windows Mobile. Implementations exist also for other operating systems.

The main advantage of ISATAP is that hosts using ISATAP have a globally unique IPv6 address even if their IPv4 address is a private one. However, there is the drawback of introducing a single point of failure into the network because the ISATAP gateway serves all IPv6 traffic.

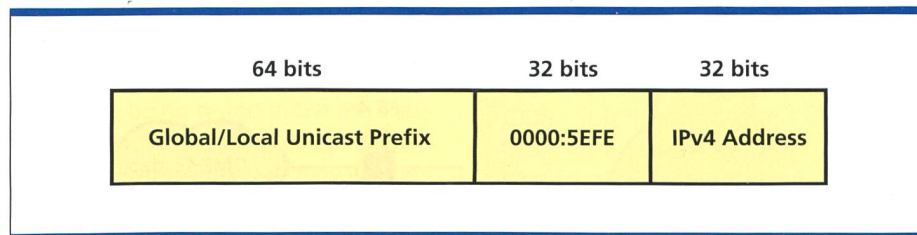


Fig. 2. ISATAP encapsulation of the IPv4 in the IPv6 address.

Additionally, performance will become a problem once a larger number of hosts use just one ISATAP gateway.

Swisscom Innovations Intranet Transition to IPv6

When planning for IPv6 introduction in our operational Intranet, the following limitations and requirements had to be addressed:

- Some of the routers in place were rather old and not upgradeable to IPv6.
- The firewall systems did not or only partly support IPv6.
- Our operations department did not yet trust IPv6 security and stability and required to have a quick and simple procedure to disconnect IPv6 traffic in case of problems.

These issues made a direct transition to a fully dual-stacked network near impossible. Having in mind the ease-of-use as the most important target, automatic tunnelling using ISATAP was the option to go for as a first step in the transition. This first phase would serve in gaining experience and in building the much needed trust. As a second step, small dual-stack islands will be added. The growing population of dual-stack islands will finally lead to the phasing out of the ISATAP solution.

Pointers

The IETF IPv6 Operations (v6ops) Working Group is developing transition scenarios and assessing suitability of transition mechanisms for different networks:

<http://www.ietf.org/html.charters/v6ops-charter.html>

The Swiss IPv6 Task Force aims at raising IPv6 knowledge and awareness and functions as a catalyst for IPv6 deployment in Switzerland:

<http://www.ch.ipv6tf.org/>

Network Design

Figure 3 shows the logical network design of the Intranet of Swisscom Innovations. A configured tunnel to Swisscom Enterprise Solutions (IP-Plus) serves as an IPv6 uplink. A separate IPv6 access router is used besides the existing IPv4 access router. Behind the IPv6 access router, native IPv6 links to both the operational and to the test DMZ were established. Each DMZ consists of two IPv6 firewalls. The operational Intranet has now two separate domains: one native IPv4 and one native IPv6. The two domains are linked by a dual-stack router serving as ISATAP gateway.

Setup and Operations at Server Side

IPv6 support varies considerably amongst different Linux distributions. Currently, the best support can be found on Red Hat Linux. We use an x86 server with Red Hat Linux 8 and kernel 2.4.20 for the ISATAP gateway. The stock Linux kernel does not support ISATAP out of the box. The USAGI project [5] provides a package with the complete kernel source tree that contains the IPv6 stack for Linux (with ISATAP support), the required libraries and userland utilities. Compilation and installation of this kernel with the additional software are consistent with regular Linux kernel compilation. IPv6 and ISATAP support have to be enabled in the kernel configuration.

Once the new kernel is installed, ISATAP support can be enabled manually using the ip tool from the USAGI distribution. For better integration with the configuration scripts of Red Hat Linux, scripts are programmed to allow automatic configuration. The scripts also start a routing advertisement demon on this interface, so that attached ISATAP nodes can query the global prefix.

This ISATAP router has now been running flawlessly for several months. The service is still only used by very few users, because its availability has not yet been widely communicated. As a conse-

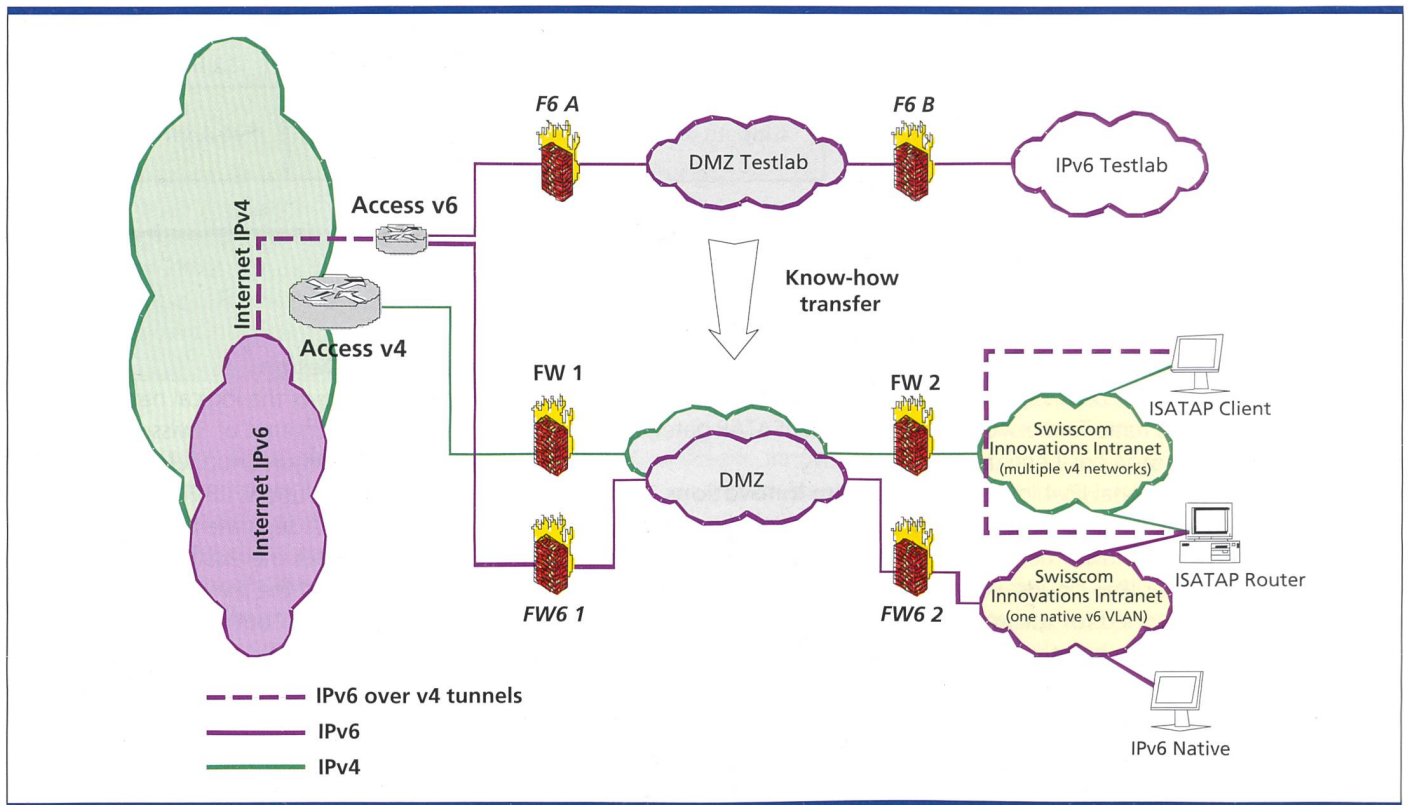


Fig. 3. Logical IPv6 network design at Swisscom Innovations.

quence, we experienced no performance bottlenecks or even overloads. The ISATAP router is periodically supervised to check its status, using tools from the Open Source domain, such as ttt [6] for real-time traffic monitoring and Nagios [7] for status surveillance. The support for IPv6 in commercial network management software is still in its infancy. The IPv6 management information, for commercial as well as for Open Source software, is mostly transported over the IPv4 network. Therefore, single-stack IPv6 networks are quite difficult to manage today. During the last few months, most of the network management tool vendors have started to IPv6-enable their tools and applications. We will include these tools in our IPv6 environment as they become available. The two firewalls at the borders of the DMZ of the IPv6 network are two servers running OpenBSD, an Open Source variant of BSD Unix. OpenBSD is renowned as a highly secure operating system, and currently it sports the best Open Source firewall for IPv6 (stateful filtering for TCP, UDP and ICMP traffic). Interested readers should refer to the OpenBSD website [8] for documentation on the operating system and the firewall configuration. Currently, commercial firewall manufac-

turers are updating their products to IPv6. Checkpoint, as an example, is already delivering quite advanced products. Other vendors like Cisco and NetScreen recently announced their commitment to support IPv6 in their security products.

Setup and Operations at Client Side

Installation of ISATAP on Windows XP with Service Pack 1 is straightforward since it is automatically activated with the IPv6 network stack. The command "ipv6 install" on the Windows XP command line activates the already built-in network stack. This process was documented on an internal website and tried out by users who fulfilled the software requirements. We have had no reports of unsuccessful configuration apart from a general incompatibility between the IPv6 network stack and a program that heavily interfered with the IPv4 network stack. The network performance experienced was on par with normal IPv6 access. Most, if not all, of our users run ISATAP on their Windows XP PCs. Having the basic IPv6 connection set up they start asking what they can do with it. Unfortunately, there are still few applications running on top of IPv6. Most users just

browse IPv6-enabled websites with the Internet Explorer. Since our web proxies are not yet IPv6-enabled, the users have to manually configure their web browser to not use the web proxy when accessing sites in the IPv6 Internet. Other software used are ThreeDegrees [9] as peer-to-peer application, Windows Media Player for multimedia streaming, and SSH for remote access to other IPv6 machines.

General Issues

Apart from software specific issues described above, one general issue related to hardware was observed. It is based on the fact that IPv4 makes little use of multicasting whilst IPv6 makes extensive use of it with the neighbour discovery protocol, the equivalent to IPv4 ARP. Network cards or their drivers designed for legacy IPv4 do not expect the Internet Protocol to rely on multicasting. Some of them (fortunately just a few) do not support that. A way around this problem is to put the cards into promiscuous mode, so that they accept all packets from the Ethernet. However, this has a negative impact on performance.

Conclusions

The experiences with IPv6 deployment in an enterprise network show that today it

is feasible for early adopters to implement IPv6. Basic network services are stable. Furthermore, ISATAP is a suitable mechanism for IPv6 deployment in an enterprise environment as it requires no upgrades in the IPv4 network and its usability is very good.

Direct migration to full dual-stack networks is likely to fail. One reason is of technical and operational nature such as limited support available with firewalls and management tools (hard factors). But another reason is the lack of knowledge and trust in IPv6 stability in the operations departments (soft factors). Further development is therefore needed not just in the technical aspects but also in education and raising awareness.

Outlook

It will be interesting to observe how the ISATAP solution performs once a big number of clients is (heavily) using the

service. As a next step, dual-stack islands will be introduced in some stub subnets. In parallel, more network services are continually being tested in the IPv6 labs of Swisscom Innovations. Once they prove to be stable enough they will be introduced in the operational Intranet. The final goal is to fully enable IPv6 in all parts of the Intranet. However, this is a long process and will still take many years.

As IPv6 brings a much simplified network architecture, application developers will make use of that. End users with access to IPv6 networks can look forward to a broad range of more secure, mobility-supporting, peer-to-peer applications easing direct collaboration and rich person-to-person communication over network domains. This is a very promising prospect and keeps the IPv6 community working towards full deployment.

6

Abbreviations

| | |
|--------|-------------------------------------------------|
| ARP | Address Resolution Protocol |
| CIDR | Classless Inter-Domain Routing |
| DMZ | De-Militarised Zone |
| DNS | Domain Name Service |
| DoD | Department of Defense |
| IETF | Internet Engineering Task Force |
| IP | Internet Protocol |
| ISATAP | Intra-Site Automatic Tunnel Addressing Protocol |
| ISP | Internet Service Provider |
| NAT | Network Address Translation |
| SSH | Secure SHell |
| USAGI | UniverSAl playGround for Ipv6 |

References

- [1] DoD memorandum on transition to IPv6: www.dod.mil/news/Jun2003/d20030609nii.pdf, July 2003.
- [2] Silvia Hagen: "Die Internet-Evolution", Comtec, April 2003.
- [3] B. Carpenter, K. Moore: "Connection of IPv6 Domains via IPv4 Clouds", IETF RFC 3056, February 2001.
- [4] F. Templin et al.: "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", IETF Internet Draft <draft-ietf-ngtrans-isatap-13.txt>, March 2003.
- [5] USAGI project: www.linux-ipv6.org, July 2003.
- [6] ttt website: www.csl.sony.co.jp/person/kjc/software.html#ttt, July 2003.
- [7] Nagios website: www.nagios.org, July 2003.
- [8] OpenBSD website: www.openbsd.net, July 2003.
- [9] ThreeDegrees website: www.threedegrees.com, July 2003.

Zusammenfassung

Das neue Internet Protokoll Version 6 (IPv6) ist eine Voraussetzung für ein nachhaltiges Wachstum des Internets. Die Wiederherstellung der End-zu-End-Transparenz ermöglicht aber auch einen Innovationsschub bei Internet-Applikationen. IPv6 ist nicht automatisch abwärtskompatibel zum heute benutzten IPv4. Um dennoch eine Koexistenz von IPv4 und IPv6 zu erlauben, wurden verschiedene Transitionsmechanismen entwickelt. Einer dieser Mechanismen, ISATAP, ist für den Einsatz in Firmennetzen gedacht. Swisscom Innovations hat nach mehrjährigen Labortests mit IPv6 nun das gesamte operationelle Intranet mit ISATAP IPv6-fähig gemacht. Die Erfahrungen zeigen, dass heute IPv6 durchaus in Firmennetzen einsetzbar ist. Allerdings muss eine Transition in Phasen durchgeführt werden. Ein direkter Schritt zu einem Dual-Stack-Netz ist nicht nur wegen der Aufrüstung der Infrastruktur (IPv6-Support in Firewalls, Network Management Tools, Router) schwierig sondern auch wegen der Schulung und Sensibilisierung des Personals. Diese ist nötig, um Vertrauen in die neue Technologie zu schaffen.

Andreas Schmid received his diploma in electrical engineering from the ETH Zürich in 1998 performing his master thesis at IBM research in Rüschlikon. Additionally, he received an Executive MBA in business engineering from the University of St. Gallen (HSG) in 2003. He has been working as an engineer and project manager for Swisscom Innovations for 5 years, interrupted by a half-year stay at the former Swisscom daughter tesion Telekommunikation in Stuttgart, Germany. His interest lies in various aspects of data networks and networking business. He is one of the founders, driving forces and currently the chairman of the Swiss IPv6 Task Force.

Martin Gysi received his masters degree in communication systems from the Swiss Federal Institute of Technology (EPFL, Lausanne) and Institut Eurecom (Sophia-Antipolis, France) in 2001, performing his thesis at the NEC European Research Laboratories (Heidelberg, Germany). Since then he has been working at Swisscom Innovations in the area of wireless and wireline access technologies (WLAN, xDSL), and future networks (IPv6).

Michael Schädler studied electrical engineering at the Swiss Federal Institute of Technology (EPFL, Lausanne) and obtained his masters degree in 2001. Since then he has been working at Swisscom Innovations in the area of wireless access technologies (WLAN, Bluetooth) and future networks (IPv6).