Deployment of Service Aware Access Networks through IPv6

Tim Stevens, Koert Vlaeminck, Wim Van de Meerssche, Filip De Turck, Bart Dhoedt, Piet Demeester
Department of Information Technology, Ghent University - IMEC
Sint-Pietersnieuwstraat 41, B-9000 Gent, Belgium
Email: tim.stevens@intec.ugent.be, koert.vlaeminck@intec.ugent.be

Abstract— Substantial research effort is currently devoted to the development of a future, low-cost access and edge network, enabling the delivery of broadband multimedia services to subscribers. As the world is moving away from current ATM (Asynchronous Transfer Mode) and PPP (Point-to-Point Protocol) access network architectures for a number of reasons, including lack of Quality of service (QoS) and autoconfiguration support, this paper presents how IPv6 can fill the hole that ATM and PPP leave behind in a very convenient, packet-oriented way. The described architecture allows multiple edge nodes to be present, increasing robustness in case of edge node failure, but also increasing scalability as non-local traffic can be balanced across multiple edge nodes. The ability to process local traffic without edge involvement further increases scalability, due to the growing peer-to-peer traffic volume. Since the architecture is independent of the underlying layer-2 technology used for the aggregation network, a phased and smooth migration from ATM to Ethernet equipment, which is claimed cheaper and easier to maintain, is possible. For some selected broadband access topologies, an efficient IPv6 subnetting scheme, minimizing routing tables through hierarchical addressing, is introduced and evaluated. Solutions for forcing non-local traffic to the correct ISP edge, in a multi-ISP environment, are presented. Finally, the deployment of the described architecture in an actual IPv6 test bed is discussed.

I. INTRODUCTION

Several shortcomings of operational DSL access networks prevent further generalization of the Internet and ubiquitous network connectivity. The current connection-oriented approach has been identified as a limiting factor both in terms of access network scalability and subscriber terminal autoconfiguration, due to the PPP (Point-to-Point Protocol) connection setup. Furthermore, a novel access network solution will have to support several services, all imposing their own QoS (Quality of service) requirements and obsoleting the single best-effort access principle as it exists today.

On the other hand, providing IP connectivity for an increased number of services, subscribers and devices per subscriber line will exhaust the available IPv4 address space faster than expected, necessitating the production deployment of IPv6. Besides the enormous address space additional IPv6 benefits for access environments include advanced QoS and multicast support, stateless autoconfiguration and standard security features, which were not present in IPv4.

This paper presents a connectionless IPv6-based data plane solution for access and aggregation networks. Given the decreasing cost of routing hardware and the need for service awareness in the aggregation nodes, the opportunity—and need—arises to upgrade access nodes to the network layer. Although some minor routing issues were identified, hierarchical address assignment, encouraged by current IPv6 policies, shows to be a perfect fit for tree-like aggregation topologies. Incorporating network-layer redundancy and processing NAP (Network Access Provider) -local traffic without ISP (Internet Service Provider) edge involvement in the hierarchical design can be obtained in a straightforward way.

Before disclosing configuration details of a test bed validating the architecture presented in this paper, the non-trivial task of migrating towards the connectionless IPv6 architecture is briefly addressed.

The paper is organized as follows. First, related work is referenced in Section II. Section III motivates the need for an IPv6 data plane for access networks. While Section IV discusses IPv6 hierarchical addressing and allocation efficiency considerations in general, Section V focuses on mapping IPv6 to tree-like aggregation structures, including IP scalability advantages such as the possibility to configure multiple edge nodes and shortest path routing. In Section VI, a brief introduction to migration aspects is presented. Before concluding the paper and identifying interesting future work in Section VIII, our test bed configuration is detailed in Section VII.

II. RELATED WORK

At present, the evolution of residential broadband connectivity is considered indispensable for further Internet generalization. Several IST projects (e.g. MUSE [1], NOBEL [2], BREAD [3]) focus on different parts of the network (i.e., access, metro and core networks) to achieve an integrated solution. While the above mentioned projects are not IPv6-specific, they clearly indicate the need for an upgraded subscriber broadband network.

Similar published work in this area mainly focuses on IPv6 deployment over existing DSL access network architectures [4], leaving existing access architectures intact. The need for an IP-aware access network for triple-play services is raised by both S. White et al. [5] and S. Sargento et al. [6]. However, both papers approach the issue from a QoS and application perspective, rather than from an IPv6 point of view.

Tim Stevens, Koert Vlaeminck, Wim Van de Meerssche, Filip De Turck, Bart Dhoedt, Piet Demeester

III. A CONVERGED IPv6 ACCESS NETWORK

A. Evolution towards connectionless access networks

Existing European DSL (Digital Subscriber Line) broadband access networks—as depicted in the upper left corner of Fig. 1—are often ATM (Asynchronous Transfer Mode) based, where end-users set up PPP connections inside ATM VCs (Virtual Circuits) to a central aggregation node, the BAS (Broadband Access Server), to access Internet services. While this architecture is adequate for current DSL broadband offerings, some serious shortcomings of PPP-based architectures jeopardize future omnipresence of high-speed Internet and large-scale introduction of advanced communication applications:

(i) All PPP links are terminated in a single device, the BAS, limiting scalability as broadband Internet penetration rates and user bandwidth demands increase.
(ii) Current PPP access networks are tailored to the connection of a single device per subscriber. Therefore, NAT (Network Address Translation) is applied on subscriber lines where multiple IP devices are connected, introducing the undesirable side-effect of breaking end-to-end IP connectivity (and transparency).
(iii) The setup of PPP links hampers device autoconfiguration (ubiquitous data link layer connectivity, IP address, default gateway, DNS (Domain Name System) server), because PPP links—carrying IP traffic—cannot be autoconfigured as the link specification is location-dependent.
(iv) PPP is conflicting with application-based QoS and multicast support.

Triggered by the current success of the Internet and IP-based corporate and home networks that are often deployed on top of an Ethernet data link layer, connectionless alternatives are considered for Internet access networks too. At present, both Ethernet access and full IP alternatives are considered as viable successors for the legacy ATM platform.

B. Ethernet versus network layer aggregation

Ethernet is evolving from a LAN (Local Area Network) solution to an all-purpose networking technology supporting first mile [7], LAN, MAN (Metropolitan Area Network) and WAN (Wide Area Network) environments [8] [9]. As a logical extension to this evolution, access network researchers are investigating Ethernet too, to replace ATM in access and aggregation networks. While some of the PPP problems mentioned above are solved by introducing Ethernet up to the edge, new ones are created. One of the main issues with access networks based on bridged LAN technology, is per subscriber segregation of layer-2 traffic, necessary for the NAP to vouch for subscriber privacy and security. The lack of QoS support in standard Ethernet is another obstacle that has to be resolved. While the introduction of VLANs (Virtual LANs) [10] reduces above mentioned Ethernet shortcomings, it can be questioned whether its 4094 VLAN identifiers (12 bits) and eight priority classes (3 bits) will provide sufficient traffic segregation facilities for very large-scale or future access topologies.

When access nodes terminate layer-2 on their UNI (user to network) and NNI (network to network) interfaces, issues for segregating different subscribers' layer-2 traffic are not present. Terminating layer-2 implies that Ethernet reframing will be necessary in the access nodes, of course. There are two options to implement layer-2 termination in the access node:

(i) Access nodes remain layer-2 devices but maintain—non-standard—state information to *link* UNI and NNI layer-2 networks;
(ii) Access nodes are layer-3 devices, using a simple *routing* algorithm to forward UNI and NNI packets.

While option (i) looks appealing at first glance, address resolution complications arise when the IP layer is mapped on top of this non-standard Ethernet behavior: the residential gateway—connected to the access node UNI interface—will expect the edge node to reside on the same IP subnet, since there is no intermediate IP node. Considering the continuously decreasing cost of processing power, it should become feasible to deploy layer-3 access nodes with a high port density in a cost-effective way. Since the access node is the first aggregation point in the network, concentrating traffic from a limited amount of subscribers, routing logic can be kept simple and straightforward, mimicking switching functionality at layer-3.

Because an increasing number of subscribers deploy a (wireless) LAN to connect multiple devices to the Internet, ISPs should investigate the distribution of IP subnets rather than single IP addresses. Currently, NAT bypasses this issue, but in the long run end-to-end connectivity is indispensable for several applications, including remote troubleshooting or configuring of networks and terminals of non-technical subscribers. Transparency is also desired in the case of subscriber nomadism, for authentication purposes. Given the shortage of IP addresses, providing a separate subnet per subscriber line is not feasible using IPv4, however. Therefore, a simultaneous migration effort towards an IPv6 network layer would be very desirable.

C. Converged access networks

As depicted in the upper left part of Fig. 1, analog voice calls and DSL data connections both set up a switched path through their respective access networks. For PSTN (Public Switched Telephone Network) calls, the switched path is set up between both end-devices and for DSL Internet access, a path is set up between a modem at the customer premises and the BAS at the access network edge. In addition to PSTN and DSL, which are physically separate except in the first mile, many people also have cable access within their homes for television distribution and/or Internet access. This means each service (e.g. telephone, Internet, television) is provided over its own access network, with a very limited shared infrastructure base.
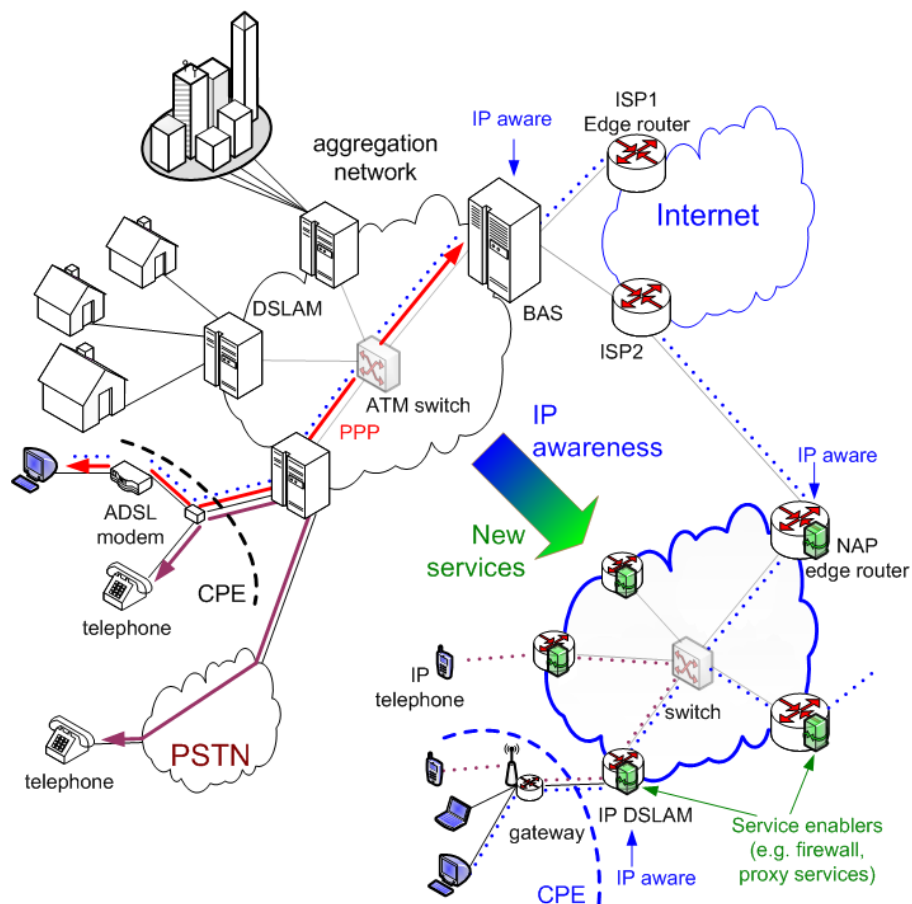
Fig. 1. Evolution towards an IP aware full service access network. The current DSL/ATM based access network is depicted in the upper left corner, with a PPP connection for each connected subscriber. Only the BAS and subscriber terminals are IP aware. Regular voice traffic is switched towards the PSTN network. In the lower right corner, a future converged IPv6 aggregation network is outlined, allowing the deployment of revenue generating services. Arrows indicate a connection is set up, while dotted lines indicate connectionless packet processing.

To reduce the capital and operational expenses of managing geographically similar, per-service separated access networks, a single converged access network as depicted in the lower right part of Fig. 1 could be deployed. Because each service has its own requirements for the characteristics of the underlying network (e.g. bandwidth, delay, jitter), advanced QoS support will be a critical success factor for future subscriber appraisal. Because Ethernet, the only serious candidate layer-2 protocol for replacing current access network infrastructure, misses standardized and mature QoS schemes as they exist for IP networks, layer-3 QoS support in access nodes appears to be the most suitable solution.

Besides the need for a larger IP address space, other data-layer requirements of novel access networks identified in this section point to the convenience of IPv6. It provides enhanced QoS and multicast support and simplifies routing by motivating hierarchical address assignment. Moreover, IPv6 enables stateless autoconfiguration, while offering the opportunity to apply stateful autoconfiguration or a combination of both, if necessary. The next section discusses IPv6-specific access network address allocation possibilities, while corresponding NAP topologies are discussed in Section V.

IV. IPv6 HIERARCHICAL ADDRESSING

A. Subnetting scheme

An IPv6 address consists of 128 bits, equally divided between the routing prefix and the interface identifier [11]. The routing prefix of a global unicast address is a globally unique prefix identifying the subnet the interface belongs to, while the last 64 bits of the IPv6 address represent the—also globally unique—EUI-64 identifier of the network interface. Global unicast subnet prefixes are assigned by the Internet Assigned Numbers Authority (IANA) through their Regional Internet Registries (RIRs). At this moment, the standard policy of RIRs (e.g. RIPE for Europe) is to assign 32-bit prefixes to ISPs [12], allowing ISPs to structure their topology with the remaining bits of the prefix. As depicted in Fig. 2, the remaining bits of the routing prefix, 32 when the default policy is applied, allow to assign prefixes to subscribers in an efficient, aggregatable way. This means the following virtual "zones" appear within the unassigned prefix bits:

(i) NAP selection bits: only necessary when an ISP operates with different NAPs within the same country, otherwise these bits can be omitted.

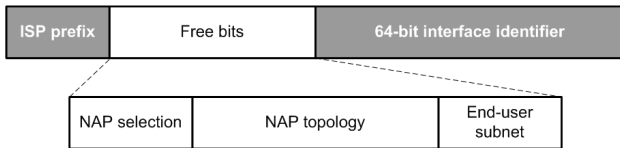(ii) Some bits representing the NAP access network topology, necessary to enable efficient routing in the

Fig. 2. IPv6 unicast address structure

access network.

(iii) Some bits representing the subscriber network topology: can be omitted if no end-user subnetting is allowed.

In this scenario, it is assumed that NAPs are responsible for the completion of the bits representing the access network topology. This includes the NAP has to configure its DSLAMs and edge routers for each ISP he is servicing, in order to advertise the correct prefixes to the subscribers. The number of bits reserved for representing the subscriber network topology may also be ISP dependent.

B. Practical considerations

Because the number of bits necessary for each part is highly dependent on the topology of the access network and the NAP and ISP policies, no fixed numbers are specified. In practice, boundaries will have to be defined based on agreements between NAPs, ISPs and subscribers. In order to avoid frequent renumbering due to future growth and network topology restructuring, some margins have to be provided, causing a certain address allocation inefficiency. This is not a new problem, as every large network (e.g. telephone networks, the current IPv4 network) sacrifices some address space to implement a manageable hierarchy. In [13], a measure of address allocation efficiency, the host density ratio (HD-ratio), is introduced as follows:

$$\mathrm{HD} = \frac{\log(\#\mathtt{allocated\_addresses})}{\log(\mathtt{max}\ \#\mathtt{available\_addresses})} \quad (1)$$

Equation (1) yields a real number between 0 and 1 for any number of allocatable addresses > 1, and any number of allocated objects ≥ 1 and $\leq$ the number of allocatable addresses.

Some case studies presented in [13] demonstrate that an HD-ratio of 80% indicates a good tradeoff between address allocation efficiency and network manageability. An HD-ratio smaller than 80% implies an inefficient address allocation policy, while networks with an HD-ratio larger than 80% quickly become unmanageable. From previous experiences with different types of networks (e.g. DECnet IV, US and French telephone network), 87% appears to be the practical upper limit at which operators are ready to make drastic decisions. Currently, the IPv4 Internet has an HD-ratio of 91% and NAT (Network Address Translation) functionality is required to maintain its manageability.

From (1), it is straightforward to calculate the practical number of bits in the IPv6 prefix necessary to represent the NAP access network topology, given the number of subscribers and a desired HD-ratio:

$$\#\mathtt{bits} = \left\lceil \frac{\log(\#\mathtt{subscribers})}{\log(2) \times \mathtt{HD}} \right\rceil, \quad (2)$$

where $\lceil x \rceil$ denotes the smallest integer larger than or equal to x. In (2), we assume each subscriber to have at least one IPv6 subnet (i.e. a 64-bit prefix). Table I summarizes results for three European countries, provided that a NAP is capable of accessing each family in a country, where a family consists on average of 2.5 persons. The calculated number of bits for a given HD ratio and population presented in Table I, represent the number of bits necessary for NAP selection and subnetting as depicted in Fig. 2. Given the current RIPE policy of distributing 32-bit prefixes per ISP, the numbers show it should be feasible to assign at least one IPv6 subnet per subscriber. While the number of bits necessary for representing the access network can be reduced by weakening the assumption each ISP should be capable of accessing each household in a country, current IAB (Internet Architecture Board) and IESG (Internet Engineering Steering Group) guidelines [14] recommend 48-bit prefix assignment for home network subscribers. This means that the current practice of assigning a 32-bit prefix per ISP—as applied by RIRs—might be too restrictive, especially for large countries. Recent RIR address allocations show ISPs can obtain larger prefixes if the assignment can be justified, however.

Because IPv6 allows one physical interface to have multiple addresses (alternating prefixes), an ISP can distribute its global prefix into the access network, without interfering with other ISPs operating in the same network (illustrated by the arrows in Fig. 3). This favors the multi-ISP scenario while maintaining layer-3 structure and manageability in the access network. As a side-note, it should be mentioned that the presence of multiple ISPs potentially decreases the access network HD-ratio, especially in the case of multihoming. Applying stateful autoconfiguration can reduce this inefficiency by means of the DHCPv6 (Dynamic Host Configuration Protocol) *prefix delegation* option [15].

V. NAP IPv6 CONFIGURATION

In this section, IPv6 is mapped on a typical tree-like access network architecture, indicating its main purpose of aggregating and interconnecting a large number of relatively small subscriber networks. Because both IPv6 addressing and access network topologies exhibit a hierarchical structure, this can happen in a straightforward manner. This section also describes IP advantages over layer-2 aggregation in terms of reliability and scalability and solutions for the access network routing issue of forcing traffic through the correct ISP edge. The section concludes discussing the future role of the ISP in layer-3 access networks.

A. IPv6 mapping to tree-like access topologies

Figure 3 presents the most straightforward layer-3 architecture for the access network, assuming access

		Required number of bits		
		Theoretical min.	*Practical min.*	*Advised*
Country	*#subscribers*	*(HD-ratio: 100%)*	*(HD-ratio: 87%)*	*(HD-ratio: 80%)*
Belgium	4 million	22	26	28
France	24 million	25	29	31
Germany	32 million	25	29	32

multiplexers are IP aware on all UNI and NNI. The root-level nodes represent different ISP's edge routers, while the intermediate nodes are NAP edge routers. This architecture allows for efficient address allocation by aggregating IPv6 prefixes on both the NAP access and edge nodes. The prefix lengths indicated on the right of Fig. 3 illustrate this behavior. On the same figure, the disadvantages of this strictly hierarchical model can be identified. First, access and edge nodes are single points of failure. Introducing redundancy in the access node might not be realistic from a technical and economical perspective, but the single edge node should definitely be avoided, as this has been identified as one of the main reasons for upgrading the current ATM network.

As illustrated in Fig. 4, providing NAP edge router redundancy can be obtained by attaching two or more edge routers to the same section of the aggregation network, where each router announces the same subnet prefixes. Using IPv6, access multiplexers do not need to run a dynamic routing protocol (e.g. Open Shortest Path First) to cope with edge router failures: IP datagrams can be forwarded based on IPv6 anycast features. IPv6 anycast addresses are indistinguishable from global unicast addresses, but apply to a group of interfaces. The interface with the shortest path to the source fulfills the request. While redundancy can be provided in this way, anycast addresses do not allow to dynamically balance the load between the redundant routers. Running dynamic routing protocols on the access multiplexers could provide additional load-balancing facilities (e.g. by forwarding traffic from/to different ISPs to different NAP edge routers). When access nodes are layer-2 devices however, the subscriber is forced to configure his own default routing policy that should be able to cope with edge router failures. This obviously limits autoconfiguration possibilities and reduces access network transparency for the end-user. At the same time, the NAP's traffic engineering possibilities are restricted, resulting in an undesirable situation.

Thus far, possible locality of traffic (i.e. data sent and received within the same access network) has not been taken into account. While this is not strictly necessary from a functional point of view and might even be considered as undesirable by ISPs, it could avoid bottlenecks in the access networks due to the increasing peer-to-peer traffic volume, currently dominating other Internet traffic [16]. As depicted by the dotted arrows in Fig. 5, local traffic is offloaded from edge nodes. Again, offering these features within a layer-2 access network is unnatural and forces the layer-2 aggregation device to perform layer-3 operations.

B. Access and edge node addressing

Several IPv6 ICMP features, including MTU path discovery, require all nodes on an IPv6 path to be globally reachable. This implies access and edge nodes need global unicast addresses even for their NNI interfaces. Since NAP nodes are ISP-independent, their NNI addresses should be allocated from a separate address pool too, possibly obtained from an IX (Internet Exchange) connecting the NAP nodes to the Internet, as depicted in Figure 6. This way, a separate—physical or virtual—path between the access nodes and the Internet, labeled "control" on the figure, is created, independent of subscriber ISPs.

C. Access network routing issue

Traffic destined to a foreign access network, needs to be routed through an ISP edge router. Because IP routing tends to follow the shortest path, it is possible a
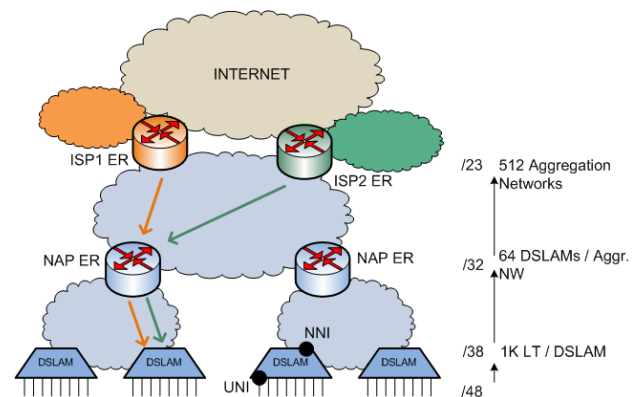

Fig. 3. Basic tree topology for routed access networks
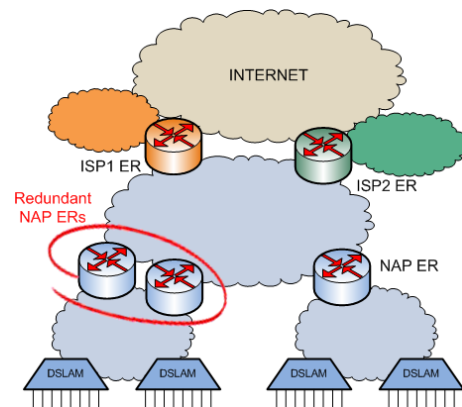

Fig. 4. Connectionless access networks allow for a seamless deployment of redundancy measures in aggregation networks
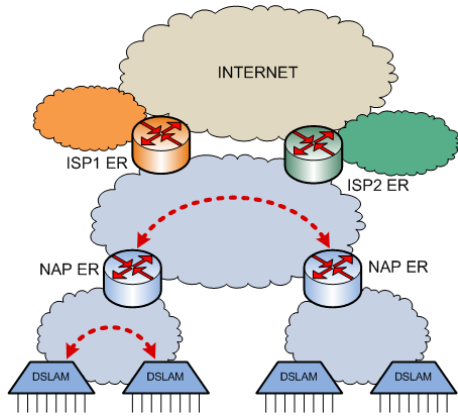
Fig. 5. Edge router interconnection to handle NAP-local traffic without ISP edge involvement

NAP edge router decides to forward data from one ISP's subscriber to a competitor ISP edge router because this router advertises a shorter route to the destination. Of course, ISPs will not forward data that does not generate any revenues. Three solutions were identified to cope with this issue:

 (i) *Policy-based routing*: access network nodes select the correct ISP edge router based on the source address.
 (ii) *IPv6 routing header*: an IPv6 routing header [17] can be attached to each IP datagram destined to another NAP.
 (iii) *IPv6-in-IPv6 tunnelling*: a NAP can set up IPv6 tunnels in the access network, forcing non-local traffic to the correct ISP edge.

Policy-based routing obviously complicates the access network routing logic and limits the applicability of common routing algorithms, but has the advantage it does not create overhead in the data plane. The other two solutions allow to employ standard routing mechanisms throughout the access network, but require additional functionality in the access and edge nodes. Additionally, both solutions create overhead in the data plane (the routing header or an additional IP header) and the injection of the routing header by an intermediate node requires non-standard
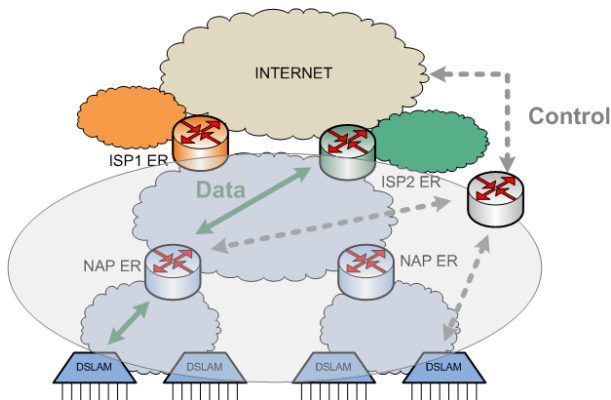
IPv6 behavior.

D. Future role of the ISP

From a technical point of view, the Network Service Provider role of an ISP is obsoleted when layer-3 access networks arise. As indicated by the previous subsection, multiple ISPs with their own address range present in the same access network even complicate the overall routing behavior. Therefore, future ISPs might shift their focus towards offering higher layer services to subscribers and evolve to Application Service Providers (ASPs). As depicted in Fig. 7, NAPs could provide network layer connectivity by assigning ISP-independent global unicast addresses to subscribers. Being a gateway between the access network and the Internet, ISPs (or IXs) can then offer global Internet connectivity as a service.

VI. MIGRATION ASPECTS

For an incumbent DSL operator, a cost-effective and seamless technical migration from the current PPP infrastructure towards a next-generation service-aware access network will be a factor of major importance. By upgrading access nodes to the IP layer, legacy ATM aggregation equipment can be used to interconnect the IP access and edge node, reducing initial investment costs and hence, risks. Figure 8 shows how ATM VCs can be constructed to provide IP connectivity between access and edge nodes, requiring no physical changes to the aggregation network. Legacy PPP connections via traditional DSLAMs—probably carrying IPv4 packets—can coexist with the new IPv6 data sent over the same aggregation network.

Since Ethernet is adopted as a solution for both Internet core networking and residential networks, migrating to Ethernet in the aggregation network—being the link between both networks—seems inevitable in the long run. These huge investment costs for operators could be spread by integrating the IP layer in access nodes, however. At the same time, this would be the ideal opportunity to trigger the wide-scale IPv6 breakthrough, a cornerstone for constructing service-aware access networks.

VII. TEST BED CONFIGURATION

For a functional evaluation of the data plane architecture presented in this paper, a small scale access network

Fig. 6. ISP-independent global unicast addresses for access and aggregation NNI interfaces
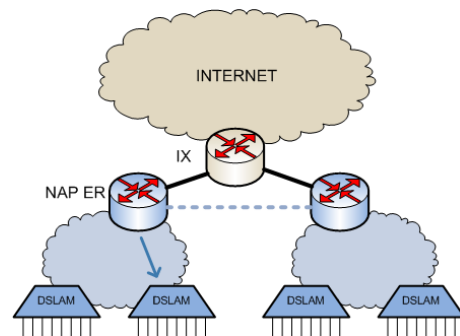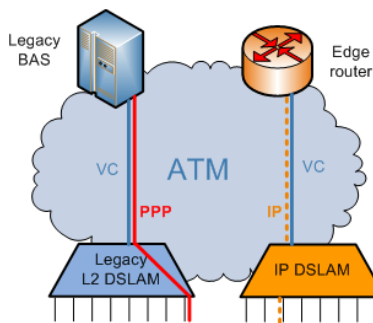

Fig. 7. ISP-less access network scenario

Fig. 8. Low-risk migration using legacy aggregation equipment

has been constructed using common PC hardware running Linux.

The aggregation network consists of eight Linux devices: three access nodes, three NAP edge routers and two ISP edge routers. They are connected in the same way as shown on Fig. 4. I.e., two access nodes are connected to a redundant pair of NAP edge routers and one access node is connected to a non-redundant NAP edge router. The two ISP edge routers present two distinct gateways to access network devices outside the NAP network.

As expected, routing issues as described in Subsection V-C occur in our dual-ISP environment. Therefore, we applied source-based routing in the access nodes, followed by either IPv6-in-IPv6 tunneling or adapting the IPv6 routing header (both are implemented). After modifying the packet, if necessary, normal routing is performed to select the correct egress interface. Source-based routing throughout all routing components in the access network would have been feasible from a technical point of view, but constrains the deployment of common IGP (Interior Gateway Protocol) routing protocols such as OSPF (Open Shortest Path First). Therefore, we decided to discard this option for our demonstrator purposes as we deployed OSPF to autoconfigure routing tables of access and edge nodes.

On the access nodes the Click modular router [18] [19] has been configured to perform non-standard operations such as IPv6 routing header manipulation. A current drawback of Click is the lack of *NIC polling mode* support for Linux 2.6 kernels, necessitating the use of 2.4 kernels for the access nodes. Linux standard interrupt-based device handling generates substantial overhead, which is avoided by the polling mechanism implemented by Kohler [19].

In order to obtain optimal Linux IPv6 support in all aggregation nodes, the USAGI [20] patch—containing an up-to-date IPv6 stack—has been applied to the standard 2.4 and 2.6 kernels.

For testing purposes, a multi-protocol server has been attached to both ISP gateways. On the free Ethernet interfaces of the access nodes, subscriber devices can be attached. The terminals connected to this demonstrator run Microsoft Windows® XP SP2 and a custom application that allows ISP selection. In fact, all ISP prefixes are configured on the host, but the application adapts the *prefix policy table* [21] to manipulate default source address selection.

VIII. CONCLUSIONS AND FUTURE WORK

As opportunities for offering new revenue-generating services in access and aggregation networks are being explored, serious shortcomings of existing connection-oriented architectures come to light. The current single edge architecture does not scale as the number of subscribers and per-user bandwidth demand increase and envisioned prospective communicational services are often based on P2P connectivity or multicast, which are not handled efficiently in existing access networks. Furthermore, offering multiple services on a physically converged access network would require application-based QoS support in the access environment, implying layer-3+ awareness in access and edge nodes. Additionally, the stringent need arises for subscriber device auto- or zero-configuration possibilities, in order to achieve ubiquitous network connectivity.

This paper shows how an IPv6 data-plane could be used as the cornerstone of a future service-oriented access network, providing tools to cope with the issues mentioned above. Additionally, IPv6 hierarchical addressing appears to be a natural match for tree-like aggregation networks. In Section VII, configuration details of the Linux/Click test bed illustrating the novel architecture, are provided.

For migration purposes, home network support for legacy IPv4 appliances cannot be ignored, as the complete migration effort will take several years. Within the residential gateway, an interface for connecting IPv4 equipment and tunnelling IPv4 data over the IPv6 aggregation network, can be foreseen. A complete IPv4/IPv6 dual stack access and aggregation environment to avoid this tunnelling issue is not a solution for the architecture presented in this paper, because the architecture relies on hierarchical address assignment, not feasible with IPv4 due to its limited addressing space.

Besides the data plane, operational access networks need an accurate Authentication, Authorization and Accounting (AAA) scheme. In a connection-oriented access network as deployed today a centralized AAA architecture is used. Connectionless access environments as proposed in this paper potentially need a decentralized AAA approach, significantly complicating the architecture. Standard IPv6 security features could be employed to enable a distributed authentication scheme. Furthermore, a converged access network needs an adequate end-to-end QoS architecture, for which IPv6 can provide the necessary building blocks. Both topics are considered as interrelated future work, necessary to constitute a complete access network architecture picture.

ACKNOWLEDGMENT

REFERENCES

[1] Multi Service Access Everywhere. [Online]. Available: http://www.ist-muse.org/

[2] Next Generation Optical Networks for Broadband European Leadership. [Online]. Available: http://www.ist-nobel.org/

[3] Broadband in Europe for All: a Multi-Disciplinary Approach. [Online]. Available: http://www.ist-bread.org/

[4] P. Cocquet, "IPv6 on DSL: The Best Way to Develop Always-on Services," *Proc. IEEE*, vol. 92, no. 9, pp. 1400–1407, Sept. 2004.

[5] S. White, R. Hernandez, A. Bodzinga, and G.-J. Bocker, "The Intelligent Broadband Access Network," in *Proc. Networks 2004: 11th International Telecommunications Network Strategy and Planning Symposium*, Vienna, Austria, June 2004, pp. 417 – 422.

[6] S. Sargentoa, R. Valadas, J. Goncalves, and H. Sousa, "IP-based Access Networks for Broadband Multimedia Services," *IEEE Commun. Mag.*, vol. 41, no. 2, pp. 146–154, Feb. 2003.

[7] *802.3ah, Ethernet in the First Mile*, IEEE Std., 2004. [Online]. Available: http://www.ieee802.org/3/efm/

[8] *IEEE 802.1ad: Provider Bridges*, IEEE Draft, Rev. 3.1, 2004. [Online]. Available: http://www.ieee802.org/1/pages/802.1ad.html

[9] Metro Ethernet Forum. [Online]. Available: http://www.metroethernetforum.org/

[10] *802.1Q, Standards for Local and Metropolitan Area Networks*, IEEE Standard, Rev. 2003 Edition, 2003.

[11] R. Hinden and S. Deering, "RFC 3513: Internet Protocol Version 6 (IPv6) Addressing Architecture," IETF, 2003.

[12] *IPv6 Address Allocation and Assignment Policy*, APNIC, ARIN, RIPE NCC Std. RIPE-267, 2003.

[13] A. Durand and C. Huitema, "RFC 3194: The Host-Density Ratio for Address Assignment Efficiency: An Update on the H Ratio," IETF, 2001.

[14] IAB and IESG, "RFC 3177: IAB/IESG Recommendations on IPv6 Address Allocations to Sites," IETF, 2001.

[15] R. Droms, J. Bound, B. Volz, T. Lemon, C. Perkins, and M. Carney, "RFC 3315: Dynamic Host Configuration Protocol for IPv6 (DHCPv6)," IETF, 2003.

[16] "Approaches to Controlling Peer-to-Peer Traffic: A Technical Analysis," P-Cube Inc, 2003.

[17] S. Deering and R. Hinden, "RFC 2460: Internet Protocol, Version 6 (IPv6) Specification," IETF, 1998.

[18] The Click Modular Router Project. [Online]. Available: http://www.pdos.lcs.mit.edu/click/

[19] E. Kohler, "The Click Modular Router," Ph.D. dissertation, Massachusetts Institute of Technology, 2001. [Online]. Available: http://www.pdos.lcs.mit.edu/papers/click:kohler-phd/thesis.pdf

[20] USAGI Project – Linux IPv6 Development Project. [Online]. Available: http://www.linux-ipv6.org/

[21] R. Draves, "RFC 3484: Default Address Selection for Internet Protocol version 6 (IPv6)," IETF, 2003.