HIPnet: A self-configuring multi-router Home IP network.

Chris Grundemann
CableLabs

Abstract

There are many new pressures and requirements emerging in today's home networks: The need for separation of visiting guest users from home users, community Wi-Fi services, smart grid, home automation & security, and an ever increasing number and type of IP enabled devices in the subscriber home are all strong motivations for additional routers and multiple LANs. The emergence of heterogeneous link layer technologies, machine to machine communication, IP & multicast video streaming, video content sharing, telecommuting and corporate IT requirements, and the possibility of home network multi-homing are all also driving additional complexity and new requirements into home networks.

This paper presents a novel approach to home router architecture, which applies many of the tools and protocols within the IPv6 framework in new ways in order to enable a completely self-configuring dual-stack (IPv4 & IPv6) multi-router home network capable of supporting the full range of in-home IP services. While many in this field are focusing on routing protocols and other complex, long-term solutions, the HIPnet approach leverages the existing Neighbor Discovery (ND) and DHCPv6 protocols, making it simpler and cheaper to implement in the near term while being robust enough to work for the long-term as well.

The paper explains the idea of directionless home routers, which have no hard-set LAN or WAN ports but rather use our "up detection" mechanism to elect a WAN port from the available physical interfaces in a deterministic manner. It describes how this method of up detection is able to create a logically hierarchical network even in a completely arbitrary and loop-filled physical topology, without introducing any new protocols. This paper then introduces "recursive DHCP Prefix Delegation (PD)" and an algorithm for IPv6 prefix sub-delegation, where the prefix delivered to the home router is divided into smaller sub-prefixes that are then distributed to directly connected downstream routers. It also explains a method for using bits from those delegated IPv6 prefixes to seed unique IPv4 prefixes without the need for IPv4 prefix delegation in DHCP. Finally, the paper describes hierarchical routing and how this entire system works as a whole to enable ISP failover and limited multihoming.

INTRODUCTION

Yesterday's home network

Home networks of the past have largely fit the same basic model: One home router connecting one Internet Service Provider (ISP) to one in-home Local Area Network (LAN). The IPv4 home LAN typically uses [RFC1918] "private" IPv4 address space to number networked devices. These "private" addresses are then rewritten in the IP header using Network Address Translation (NAT) overloading at the home router to allow a single "public" (globally routable) IPv4 address to be shared among all IP enabled devices in the home netowork for Internet connectivity.

More recently, many home users have started to add additional routers to their home networks, often unintentionally. It is becoming common to find home users who has purchased and deployed a second or third home router in order to extend Wi-Fi range or provide additional Ethernet ports. Although there are devices better suited to these tasks (Access Points (APs) and Repeators for Wi-Fi range and Ethernet Switches for physical ports), these “specialty” devices are less common in retail stores and often cost as much or more than the general purpose routers. Additionally, many people (consumers and retail store clerks alike) are more familiar with home routers than they are other home networking gear. This combination of familiarity, availability, and affordability seems to be driving more and more home users to (often inadvertently) deploy multi-router home network topologies.

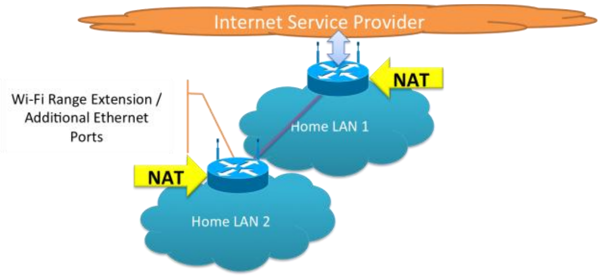
Internet Service Provider
NAT
Wi-Fi Range Extension / Additional Ethernet Ports
Home LAN 1
NAT
Home LAN 2

Yesterday’s Home Network

In the legacy paradigm of IPv4-only and NAT, this deployment of multi-router home networks is problematic. Because these routers are designed with the one-router/one-LAN architecture in mind, there are several problems introduced when using more than one of them in a home network.

First, adding an additional router creates an additional LAN, with it’s own DHCP server, address pool (which often overlaps other home LANs), and default route (through the new router). As this is usually not the intention of the home user, and may not even be a known consequence, it can cause problems when setting up new devices and services or when troubleshooting existing ones. Introducing multiple, routed LANs also stops link-local traffic from reaching the entire home network, which breaks many forms of service discovery.

Second, inserting even a second legacy home router creates introduces a second NAT, this one within the home network itself. This causes traffic between the home LANs to undergo NAT and all traffic leaving the second LAN to undergo two rounds of NAT before reaching the ISP network. These in-home and multi-layer NATs are known to impair or completely break many protocols and applications (e.g. service discovery, IPsec, DNSSEC, etc.). Additional layers of routers add additional layers of NAT, worsening the problem.

Third, these NAT routers inherently introduce a stateful firewall, which exaserbates many of the issues already raised.

Perhaps even more alarming is the state of IPv6 in the home. Unlike IPv4 NAT which allows multiple routers to be linked one behind the other to offer at least some connectivity to connected devices, current IPv6 enabled home routers [6204bis] do not support any standard mechanism to facilitate such “chaining.” This means that devices connected to a second home router are likely to not have any IPv6 connectivity outside of the LAN at all.

The end result is that multi-router home networks built with legacy home routers have limited functionality, and the functionality is further limited as size and complexity increase. There are ways to solve many of these limitations but they all require manual configuration of home routers and networked devices, which is beyond the ability of most home users.

Emerging use cases

While home IP networks have been able to remain relatively simple over their 20-30 year life so far, there are now use-cases emerging which will surely change that going forward.

One of the first trends to appear is that of "guest" networks. Many home users have an increasing amount of personal or private data on their networked devices which they may not want visitors to their home to be able to access. Family photographs, financial and tax documents, other legal materials, and many other common types of electronicly stored information are seen as sensitive. In order to provide Internet access for guests without exposing this sensitive data home users are beginning to deploy a second, "guest," SSID on their wireless networks.

Other requirements for additional home LANs include:

1) Community Wi-Fi applications where a Wi-Fi gateway in the user's home is used to provide Wi-Fi roaming services for subscribers other than the home user.

2) Femto cell applications in which a gateway in the subscriber home is used to provide cellular services.

3) Telecommuting and corporate IT requirements for network separation between business and personal LANs within the user's home.

4) The emergence of heterogeneous link layer technologies, such as Bluetooth, ZigBee, and Z-Wave, which require their own proprietary gateways.

5) The introduction of IP based or otherwise networked Security, Monitoring, and Automation systems and services.

In addition to these emerging use-cases for multi-LAN/multi-router home networks, there are several other trends towards more complex home networks. These primarily revolve around the growing amount of IP video and the ever increasing number of IP enabled devices in the subscriber home.

IP is becoming a prominent medium for transmitting video. Many home users are already streaming at least some portion of the video they consume over IP from the Internet. Likewise, IP is the de-facto standard for video content sharing and streaming between devices inside the home. This IP delivery of video places new burdens and requirements on home networks, which drives additional complexity.

The explosion of IP enabled devices in each home shows no signs of slowing in the forseable future. The smorgasboard of desktops, laptops, tablets, and mobile phones is being supplemented further with all manner of "smart" appliances, cameras, sensors, printers, storage devices, and surely more to come. This trend adds to those above and drives even more complexity into home networks.

The home network of tomorrow

The emerging home network use-cases outlined above, in addition with others not covered here and perhaps not even apperent yet, lead to a future in which complex, multi-router, home networks become commonplace.

A typical home network of tomorrow is likely to have some combination of multiple W-Fi SSIDs, a "guest" LAN or two, multiple function-specific LANs with their own routers or gateways (e.g. a LAN in the kitchen for smart appliances and a LAN in the living room for media devices, etc.), and a multitude of IP enabled services and devices running over it all.

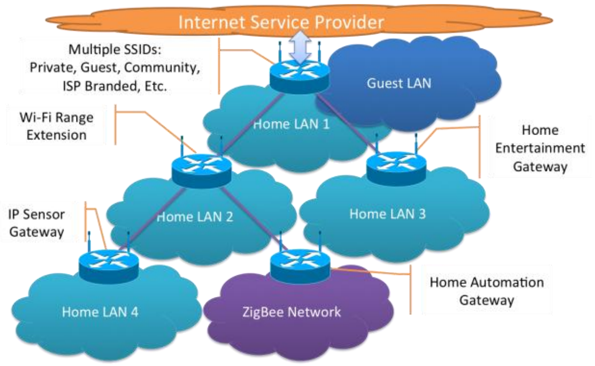

Tomorrow's Home Network

While we can not assume to predict the future with absolute certainty, we can easily see that home networks will become more complex, continuing to support more and more routers, devices, and serivces. We can also assume that these complex, multi-router home networks will need to "just work" in order to facilitate their operation by average home users.

THE HIPNET SOLUTION

HIPnet (derived from Home IP networking) is a near term solution to complex home networks. Specifically, the HIPnet solution defines a self-configuring home router architecture which:

1) Is capable of operating in increasingly large (and arbitrarily constructed) residential home networks.

2) Requires no user interaction for the vast majority of use-cases.

3) Uses existing protocols in new ways.

4) Does not require a routing protocol.

5) Meets the principles for advanced home networks defined in [homenet].

Guiding principles

Five common principles have guided the development of HIPnet:

1) *Home networks will become more complex, home users will not.* As discussed above, a multitude of emerging use-cases are driving additional complexity into home networks. Despite this, there is no reason to assume a parallel trend of home users becoming more networking savvy will emerge as well. These complex home networks of the future need to "just work" in the majority of cases for the majority of users, without any manual configuration whatsoever.

2) *Invoking a "god box" leads to religious wars.* A "god box" commonly refers to any device which attempts to be all things to all people. Here specifically a "god box" would be a home gateway intended to fulfill all home networking requirements. This is unlikely to be successful short-term in an environment of proprietary solutions nor long-term in a field that evolves as quickly as networking. Competing standards and solutions are sure to "war" with each other and make consensus on any comprehensive all-in-one solution nearly or totally impossible.

3) *New protocols bring new problems.* The introduction of any new protocol, whether completely new or simply new to that application or use, almost invariably creates new problems which must be addressed or worked around. It is almost always preferable to use the protocols already available, perhaps in slightly new ways, then to introduce the uncertainty of a new protocol.

4) *We have enough addresses.* IPv6 provides a glut of individual addresses. This allows us to be less concerned with the highest possible efficiency in address usage and to focus on simplifying network functionality.

5) *Use IPv6, support IPv4.* The IPv6 protocol suite provides us with all the tools we need to create a complete solution to complex multi-router home networks. When

creating this IPv6 based architecture, however, we mustn't forget to maintain support for legacy IPv4 devices and services.

Terminology

The following terms will be used throughout the remainder of this document to describe the HIPnet solution in detail:

Home IP Network (HIPnet) Router: A node intended for home or small-office use that forwards packets not explicitly addressed to itself.

End-User Network: One or more links attached to the HIPnet router that connect IPv6 and IPv4 hosts.

Service provider: An entity that provides access to the Internet. In this document, a service provider specifically offers Internet access using IPv6, and may also offer IPv4 Internet access. The service provider can provide such access over a variety of different transport methods such as DSL, cable, wireless, and others.

Customer Edge Router (CER): A HIPnet router that connects the end-user network to a service provider network.

Internal Router (IR): An additional HIPnet router deployed in the home or small-office network that is not attached to a service provider network. Note that this is a functional role; it is expected that there will not be a difference in hardware or software between a CER and IR, except in such cases when a CER has a dedicated non-Ethernet WAN interface (e.g. DSL/cable/ LTE modem) that would preclude it from operating as an IR.

Up interface: A HIPnet router's attachment to a link where it receives one or more IP addresses and/or prefixes. This is also the link to which the HIPnet router points its default route.

Down interface: A HIPnet router's attachment to a link in the end-user network on which it distributes addresses and/or prefixes. Examples are Ethernet (simple or bridged), 802.11 wireless, or other LAN technologies. A HIPnet router may have one or more network-layer down interfaces.

Downstream router: A router directly connected to a HIPnet router's Down Interface.

Depth: The number of layers of routers in a network. A single router network would have a depth of 1, while a router behind a router behind a router would have a depth of 3.

Width: The number of routers that can be directly subtended to an upstream router. A network with three directly attached routers behind the CER would have a width of 3.

Edge detection

Customer Edge Routers (CER) will often be required to behave differently from Internal Routers (IR) in several capacities. Some examples include: Firewall settings, IPv4 NAT, ULA generation (if supported), name services, multicast forwarding differences, and others. This is a functional role, and will not typically be differentiated by hardware/software (i.e. end users will not purchase a specific CER model of router distinct from IR models).

There are three methods that a router can use to determine if it is a CER for its given network:

1) "/48 Check" - Service providers will provide IPv6 WAN addresses (DHCPv6 IA_NA) and IPv6 prefixes (DHCPv6 IA_PD) from different pools of addresses. The largest IPv6 prefix that we can expect to be delegated

to a home router is a /48. Combining these two observations, a home router can compare the WAN address assigned to it with the prefix delegated to it to determine if it is attached directly to a service provider network. If the router is a CER, the WAN address will be from a different /48 than the prefix. If the router is an IR, the WAN address will be from the same /48 as the prefix. In this way, the router can determine if it is recieving an "external" prefix from a service provider or an "internal" prefix from another home router.

2) CER_ID - A home router can use the CER_ID DHCPv6 option defined in [CER-ID] to determine if it is a CER or an IR. ISPs will not set the CER_ID option, but the first CPE router sets its address in the option and other routers forward the completed CER_ID to subdelegated routers.

3) Physical - Some routers will have a physical differentiator built into them by design that will indicate that they are a CER. Examples include mobile routers, DSL routers, and cable eRouters. In the case of a mobile router, the presence of an active cellular connection indicates that the router is at the customer edge. Likewise, for an eRouter, the presence of an active DOCSIS® link tells the router that it is at the customer edge.

HIPnet routers can (and likely will) use more than one of the above techniques in combination to determine the edge. For example, an internal router will check for the CER_ID option, but will also use the 48 check in case its upstream router does not support CER_ID.

Directionless routers

As home networks grow in complexity and scale, it will become more common for end users to make mistakes with the physical connections between multiple routers in their home or small office. This is liklely to produce loops and improper uplink connections. While we can safely assume that home networks will become more complex over time, we cannot make the same assumption of the users of home networks. Therefor, home routers will need to mitigate these physical topology problems and create a working multi-router home network dynamically, without any end user intervention.

Legacy home routers with a physically differentiated uplink port are "directional;" they are pre-set to route from the 'LAN' or Internal ports to a single, pre-defined uplink port labeled "WAN" or "Internet". This means that an end-user can make a cabling mistake which renders the router unusable (e.g. connecting two router's uplink ports together). On the other hand, in enterprise and service provider networks, routers are "directionless;" that is to say they do not have a pre-defined 'uplink' port. While directional routers have a pre-set routing path, directionless routers are required to determine routing paths dynamically. Dynamic routing is often achieved through the implementation of a dynamic routing protocol, which all routers in a given network or network segment must support equally. This section introduces an alternative to dynamic routing protocols (such as OSPF) for creating routing paths on the fly in directionless home routers.

Note that some routers (e.g. those with a dedicated wireless/DSL/DOCSIS® WAN interface) may continue to operate as directional routers. The HIPnet mechanism described below is intended for general-purpose routers.

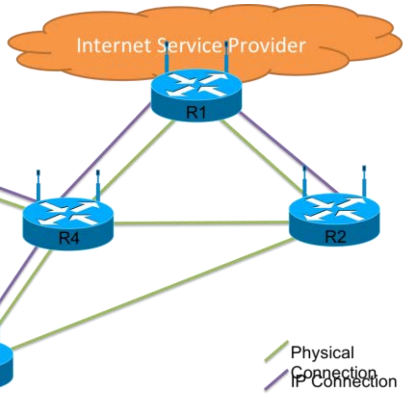

Physically Arbitrary with Logical Hierarchy

The HIPnet mechanism uses address acquisition as described in [6204bis] and various tiebreakers to determine directionality (up vs. down) and by so doing, creates a logical hierarchy (cf. [prefix-alloc]) from any arbitrary physical topology:

1) After powering on, the HIPnet router sends Router Solicitations (RS) [RFC4861] on all interfaces (except Wi-Fi*)

2) Other routers respond with Router Advertisements (RA)

3) Router adds any interface on which it receives an RA to the candidate 'up' list

4) The router initiates DHCPv6 PD on all candidate 'up' interfaces. If no RAs are received, the router generates a /48 ULA prefix.

5) The router evaluates the offers received (in order of preference):

a) Valid GUA preferred (preferred/valid lifetimes >0)

b) Internal prefix preferred over external (for failover - see below)

c) Largest prefix (e.g. /56 preferred to /60)

d) Link type/bandwidth (e.g. Ethernet vs. MoCA)

e) First response (wait 1 s after first response for additional offers)

f) Lowest numerical prefix

6) The router chooses the winning offer as its Up Interface.

Once directionality is established, the router continues to listen for RAs on all interfaces but doesn't acquire addresses on Down Interfaces. If the router initially receives only a ULA address on its Up Interface and GUA addressing becomes available on one of its Down Interfaces, it restarts the process. If the router stops receiving RAs on its Up Interface, it restarts the process.

In all cases, the router's Up Interface becomes its uplink interface; the router acts as a DHCP client on this interface. The router's remaining interfaces are Down Interfaces; it acts as a DHCP server on these interfaces. Also, per [6204bis], the router only sends RAs on Down Interfaces.

*Note: By default, Wi-Fi interfaces are considered to point "down." This requires manual configuration to enable a wireless uplink, which is preferred to avoid accidental or unwanted linking with nearby wireless networks.

Recursive prefix delegation

HIPnet routers use DHCPv6 prefix sub-delegation ([RFC3633]) to recursively build a hierarchical network ([prefix-alloc]). This approach requires no new protocols to be supported on any home routers.

Once directionality is established, the home router will acquire a WAN IPv6 address and an IPv6 prefix per [6204bis]. As HIPnet routers (other than the CER) do not know their specific location in the hierarchical network, all HIPnet routers use the same generic rules for recursive prefix delegation to facilitate route aggregation, multihoming, and IPv4 support (described below). This methodology expounds upon that previously described in [prefix-alloc].

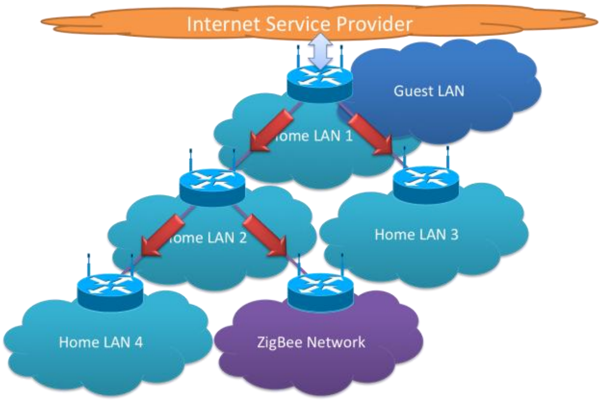

Recursive Prefix Delegation

The process can be illustrated in the following way:

1) Per [6204bis], the HIPnet router assigns a separate /64 from its delegated prefix(es) for each of its Down Interfaces in numerical order, starting from the numerically lowest.

2) If the received prefix is too small to number all Down Interfaces, the router collapses them into a single interface, assigns a single /64 to that interface, and logs an error message.

3) The HIPnet router subdivides the IPv6 prefix received via DHCPv6 ([RFC3315]) into sub-prefixes. To support a suggested depth of three routers, with as large a width as possible, it is recommended to divide the prefix on 2-, 3-, or 4-bit boundaries. If the received prefix is not large enough, it is broken into as many /64 sub-prefixes as possible and an error message is logged. By default, this document suggests that the router divide the delegated prefix based on the aggregate prefix size and the HIPnet router's number of physical Down Interfaces. This is to allow for enough prefixes to support a downstream router on each down port.

* If the received prefix is smaller than a /56 (e.g. a /60), a HIPnet router with 8 or more ports divides on 3-bit boundaries (e.g. /63) and a HIPnet router with 7 or fewer ports divides on 2-bit boundaries (e.g. /62).

* If the received prefix is a /56 or larger, a HIPnet router with 8 or more ports divides on 4-bit boundaries (e.g. /60) and a HIPnet router with 7 or fewer ports divides on 3-bit boundaries (e.g. /59).

4) The HIPnet router delegates remaining prefixes to downstream routers per [RFC3633] in reverse numerical order, starting with the numerically highest. This is to minimize the renumbering impact of enabling an inactive interface.

For example, a four port router with two LANs (two Down Interfaces) that receives 2001:db8:0:b0::/60 would start by numbering its two Down Interfaces with 2001:db8:0:b0::/64 and 2001:db8:0:b1::/64 respectively, and then begin prefix delegation by giving 2001:db8:0:bc::/62 to the first directly attached downstream router.

<u>Hierarchical routing</u>

The recursive prefix delegation method described above, coupled with "up detection", enables very simple hierarchical routing. By this we mean that each router installs a single default 'up' route and a more specific 'down' route for each prefix delegated to a downstream IR. Each of these 'down' routes simply points all packets destined to a given prefix to the WAN IP address of the router to which that prefix was delegated. This combination of a default 'up' route and more specific 'down' routes provides complete reachability within the home network with no need for any additional message exchange or routing protocol support.

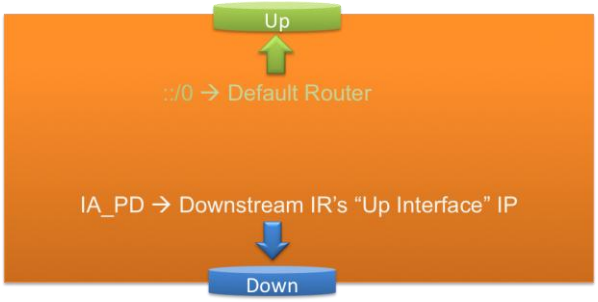

Hierarchical Routing Table

Multiple address families

The recursive prefix delegation method described above can be extended to support additional address types such as IPv4, additional GUAs, or ULAs. When the HIPnet router receives its prefix via DHCPv6 ([RFC3633]), it computes its 16-bit Link ID (bits 48-64) from the received IA_PD. It then prepends additional prefixes received in one or more IPv6 Router Advertisements ([RFC4861]) or from the DHCPv4-assigned ([RFC2131]) IPv4 network address received on the Up Interface.

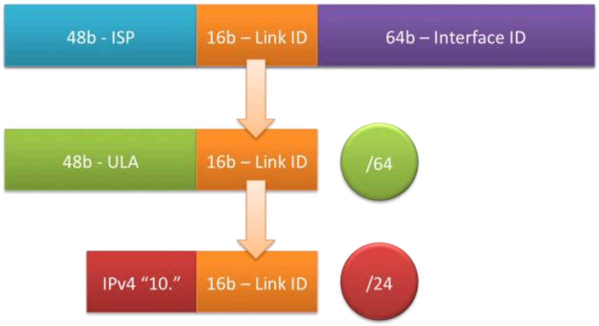

Link ID

As the network is hierarchical, upstream routers know the Link ID for each downstream router, and know the prefix(es) on each LAN segment. Accordingly, HIPnet routers automatically calculate downstream routes to all downstream routers.

In networks using this mechanism for IPv4 provisioning, it is suggested that the CER use addresses in the 10.0.0.0/8 ([RFC1918]) range for downstream interface provisioning.

Multiple ISP: Failover

Using the procedures described above, multi-router home networks with multiple ISP connections can easily operate in an active/standby manner, switching from one Internet connection to the other when the active connection fails. Lacking a default priority, HIPnet routers will have to default to a "first online" method of primary CER selection. In other words, by default, the first CER to come online becomes the primary CER and the second CER to turn on becomes the backup. In this text, the primary ISP is the ISP connected to the primary CER and the backup ISP is simply the ISP atached to the backup CER.

In an active/standby multi-ISP scenario, a backup CER sets its Up Interface to point to the primary CER, not the backup ISP. Hence, it does not acquire or advertise the backup ISP prefix. Instead, it discovers the internally advertised GUA prefix being distributed by the currently active primary CER.

In the case of a primary ISP failure, per [6204bis], the CER sends an RA advertising the preferred lifetime as 0 for the ISP-provided prefix, and its router lifetime as 0. The backup CER becomes active when it sees the primary ISP GUA prefix advertised with a preferred lifetime of 0. In the case of CER failure, if the backup CER sees the Primary CER stop sending RAs altogether, the Backup CER becomes active.

When the backup CER becomes active, it obtains and advertises its own external GUA. When advertising the GUA delegated by its ISP, the backup CER sets the valid, prefered, and router lifetimes to a value greater than 0. Other routers see this and re-determine the network topology via "up" detection, placing the new CER at the root of the new hiearchical tree.

Using this approach, manual intervention may be required to transition back to the primary ISP. This prevents flapping in the event of intermittent network failures. Another alternative is to have a user-defined priority, which would facilitate pre-emption.

Multiple ISP: Multi-homing

The HIPnet algorithm also allows for limited active/active multihoming in two cases:

1) When one ISP router is used as the primary connection and the second ISP router is used for limited connectivity e.g. for a home office.

2) When both ISP routers are connected to the same LAN segment at the top of the tree.

In case 1, the subscriber has a primary ISP connection and a secondary connection used for a limited special purpose. (e.g. for work VPN, video network, etc.). Devices connected under the secondary network router access the Internet through the secondary ISP. All devices still have access to all network resources in the home. Devices under the secondary connection can use the primary ISP if the secondary fails, but other devices do not use the secondary ISP.

As described above, the primary CER performs prefix sub-delegation to create the hierarchical tree network. The secondary edge router then obtains a second prefix from ISP2 and advertises the ISP2 prefix as part of its RA. The Secondary CER thus includes sub-prefixes from both ISPs in all IA_PD messages to downstream routers with the same "router id.". In a change from the single-homing (or backup router) case, the secondary CER points its default route to ISP2, and adds an internal /48 route to its upstream internal router (e.g. R1). Devices below the the secondary CER (e.g. Host 2, Host 3) use ISP2, but have full access to all internal devices using the ISP1 prefix (and/or ULAs). If the ISP2 link fails, the secondary CER points its default route 'up' and traffic flows to ISP1. Devices not below the secondary CER (e.g. Hosts 1, 4, 5) use ISP1, but have full access to all internal devices using the ISP1 prefix (or ULAs).

In case 2, the secondary CER is installed on the same LAN segment as the primary CER. As above, it acquires a prefix from both the CER and secondary ISP. Since it is on the same LAN segment as the CER, the secondary CER does not delegate prefixes to that interface via DHCP. However, it does generate an RA for the ISP2 prefix on the LAN.

As described above, downstream routers receiving the secondary CER RA acquire an address using SLAAC and generate a prefix for sub-delegation by prepending the secondary CER prefix with the Link ID generated during the receipt of the prefix from the CER. Such routers then generate their own RAs on downstream interfaces and include the secondary prefix as an IA_PD option in future prefix delegations.

CONCLUSION

HIPnet is a near term, scalable solution to the problem of increasing complexity in home networks. The HIPnet architecture described in this document meets the challenges of tomorrow’s home networks without the need for manual configuration or routing protocols by employing directionless home routers, recursive prefix delegation, and hierarchical routing.

While the primary focus is on IPv6 support, this document also describes how HIPnet leverages IPv6 to configure IPv4 in a manner better than the nested NATs in operation on many networks today.

This document describes how a HIPnet router automatically detects both the edge of the customer network and its upstream interface, how it subdivides an IPv6 prefix to distribute to downstream routers, and how it leverages IPv6 address assignment to distribute IPv4 addresses. It also discusses how such a router can operate with a backup ISP or limited multihoming across two ISPs.

Specific requirements for building a HIPnet compliant home router can be found in [hipnet].

REFERENCES

[prefix-alloc]   Nordmark, E., Chakrabarti, S., Krishnan, S., and W. Haddad, "Simple Approach to Prefix Distribution in Basic Home Networks", draft-chakrabarti-homenet-prefix-alloc-01 (work in progress), October 2011.

[cer-id]   Donley, C. and C. Grundemann, "Customer Edge Router Identification Option", draft-donley-dhc-cer-id-option-01 (work in progress), September 2012.

[hipnet]   Grundemann, C., Donley, C., Brzozowski, J., Howard, L., and V. Kuarsingh, “A Near Term Solution for Home IP Networking (HIPnet)”, draft-grundemann-homenet-hipnet-01 (work in progress), February 2013.

[homenet]   Chown, T., Arkko, J., Brandt, A., Troan, O., and J. Weil, "Home Networking Architecture for IPv6", draft-ietf-homenet-arch-07 (work in progress), February 2013.

[6204bis]   Singh, H., Beebee, W., Donley, C., and B. Stark, "Basic Requirements for IPv6 Customer Edge Routers", draft-ietf-v6ops-6204bis-12 (work in progress), October 2012.

[RFC 1918]   IETF RFC 1918, Address Allocation for Private Internets, Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, E. Lear, February 1996.

[RFC 2131]   IETF RFC 2131, Dynamic Host Configuration Protocol, R. Droms, March, 1997.

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

[RFC 3633]   IETF RFC 3633, IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6, O. Troan, R. Droms, December 2003.

[RFC 4861]   IETF RFC 4861, T. Narten, E. Nordmark, W. Simpson, H. Soliman, Neighbor Discovery for IP Version 6 (IPv6), September 2007.