The ISP Column
A monthly column on things Internet

July 2011
Geoff Huston

Networking @ Home

For me, one of the more interesting sessions at the recent IETF 81 meeting in July was the first meeting of the recently established Homenet Working Group.

What's so interesting about networking the home?

Well, if you regard challenges as "interesting", then just about everything is interesting when you look at networking in the home!

It's been a very long time since the state of the art in home Internet was plugging the serial port of the PC into the dialup modem. Even the ADSL modem, even when combined with some for of WiFi base station, is looking distinctly passé these days. Today the home network is seeing the intersection of a whole set of interests, including the phone service, the TV service, home security services, energy management, utility service metering, possibly other forms of home device monitoring, and, oh yes, connecting the laptops and the mobile devices to the net. And of course its not just a home LAN over a wired network. WiFi home networks are commonplace, and of course there are various Bluetooth devices. Maybe sometime soon it will be common for the home network to also host some form of 3G femtocell as well. But these days even that level of network complexity is not enough. Increasingly, the home office is part of the work office, and if there are a number of residents at home then the home network may be an endpoint for a number of corporate and institutional Virtual Private Networks (VPNs). Within all this mélange we want sophisticated security. Its not just protecting the home network from the neighbors, but the security requirements include the ability that allows individuals to partition off their work-VPN part of the home network from other home users. Oh, and for resiliency we might want a second provider, such as a mobile service to the home, so we might want to add site-based multi-homing to the mix. And now we need to make all this fly in both IPv6 and IPv4.

That's a massive agenda of requirements. But to make this situation truly challenging, we can't expect every home to come with an IT Operational Service Manager to ensure that all the various devices you bring into the home and connect to the network all function as required for the home's particular requirements. Indeed, we can't expect any home to be so lavishly supported, nor can we afford to support home networking with a bevy of specialized call centers with on-demand support specialists expert in the panoply of consumer devices that are being sold today.

With home networks the bottom line is that the consumer is effectively on their own, and all this equipment better just work straight out of the box. No configuration, no buttons, it just has work!

Routing @ Home

The evolution of networking at home has progressed from a single computer to a basic Local Area Network (LAN), and from there to an ether-bridged network with a number of WiFi and wired LAN segments. All these environments have a single common architecture of a single "boundary" unit that acts as a point of demarcation between the Internet Service Provider (ISP) and the home network. This unit is generally called Customer Premises Equipment (CPE), and typically encompasses the functions of a modem, IPv4 NAT, DHVPv4 server, DHCPv6 server, security firewall, bridge and rudimentary router.

But its unrealistic to assume that home networks will continue to use a centralized model that places the entire management functionality of the home network in a single unit. So how should we view home networks? Should home networks be a single bridged LAN, or are we seeing the evolution of home networks into multiple distinct domains with a routing fabric to glue them together? And if this is the case what routing protocol should be used?

I have noticed in the low end of the CPE market its not uncommon to see a rudimentary routing functionality supported by RIP. Now, thankfully, its RIPv2, so the routing protocol can be configured with variable length subnet masks, but even so, RIP is a very basic and simple routing protocol. But perhaps in this environment that might be a positive factor rather than a liability, in so far as RIP is simple enough to be auto-configurable. On the other hand if there is an emergent need for more complex functions then maybe we need to look a little harder at what options are available.

One of these more complex functions is the issue of subnet management. In IPv6 the CPE will collect an IPv6 address prefix. This differs from the conventional IPv4 environment where the CPE is typically assigned a single IPv4 address. So the ensuing question is: Is it possible to automate the distribution of IPv6 subnets across the entire home network? What form of management protocol is appropriate for this role.

And of course the world gets a whole lot more complicated if the home network has two (or more) service providers. In the IPv6 environment this starts to become a challenging task, not only with the distribution of multiple subnets across the home network, but also in the issue of exit path selection. If the home network is exercising due diligence to prevent source address spoofing it is also necessary for the home's routing infrastructure to deliver an outgoing packet to the "right" exit ISP, where the source address of the outgoing packet needs to match the address prefix provided by the corresponding ISP's service. In other words there is a requirement for source address routing in the home. This is a challenge that was not really addressed by the Site Multi-Homing Working Group (SHIM6), despite the best of intentions, and it represents an even greater challenge if the intent is to provide mechanisms that can achieve this in an unmanaged home network environment.

I must admit to some concern here. We've managed to keep routing work by using two principles. The first is to try and keep the routing task as simple as possible. Routing propagates a single "best" path to a destination. It does not necessarily do this quickly, nor necessarily does it carry around with it a whole set of alternatives. It does just one job. And with that we've been able to keep routing working. The second principle is to admit that we have never really succeeded with the first principle of functional simplicity and we have always had expertise at hand to oversee the routing function and apply manual patches as required! The specialized requirements for the home network appears to be breaking both principles. The requirements are certainly not simple and I see a mix of routing techniques, including various forms of policy-based routing requirements entering the discussion. Secondly, there is no assurance that if things fail there is expertise at hand to mend the failure. Indeed the more complex the routing environment the greater the potential for complex forms of failure. Indeed as we contemplate ever more complex requirements in the home network, the greater the risk of encountering failure "by design" where it is just not possible to design products for this environment that can just work.

Names @ Home

What should I call my printer? More to the point, how should I identify my WiFi printer to all those devices at home that want to use it to print. I'm sure that I would not like to use a proprietary naming scheme that requires me to add additional name resolution software to every device at home that wants to print something, nor do I want to transcribe IP addresses into everything. I'd like my printer to get dynamically assigned IPv4 and IPv6 addresses when the device is plugged in and switched on, and have the printer's name published via a generic name resolution mechanism, namely the DNS.

But most of the time the rest of the world has no need to know the name of my printer at home, and I'm not sure that it's a good move, security wise, to gratuitously publish information in the public DNS. So what I would like for my printer is some form of "local" or "scoped" DNS, where I can name my printers, my disk servers, and other devices that I have at home in the context of my home and not have this information leak further afield. Is this scoped form of name resolution, split horizon DNS, or split views, possible in the context of the DNS without invoking further elements of configuration management?

Multicast DNS (mDNS) is perhaps one of the strongest candidates for this role. In essence mDNS replaces the explicit server / client structure of the DNS with a scoped name subdomain of .local that is inherently scoped to the scope of the associated multicast domain. This allows a client to perform DNS-like name resolution functions on a local network without the need to configure a conventional DNS server environment, and without the need to obtain global delegation of a site name in the global DNS.

An alternative approach is to use a conventional DNS delegation and conventional unicast DNS queries and responses. Clients are able to use DNS Dynamic Updates to update the local DNS server with their details as they come online. (This either requires open access from anyone to the nameserver, or a security mechanism such as TSIG. TSIG generally requires manual configuration, and alternatives are either little used, such as TKEY, or start to involve further intricacies, such as Microsoft's Active Directory, which uses other user authentication mechanisms to bootstrap the TSIG part using GSSTSIG) The DNS server itself can be advertised to all clients via the Simple Service Discovery Protocol (SSDP), as part of the larger Universal Plug and Play (UPnP) framework.

Sensing and Serving @ Home

Where too from here? Its certainly the case that electronics has managed to pervade just about every device at home. Electricity meters are morphing into household energy management systems, and many other household appliances are now controlled by internal processors. But individually configuring each of these devices is a forbidding task. Even adding an interface to allow manual configuration can often be a challenging objective.

So the objective here is to define a standard mechanism to allow sensors to sense their local environment when powered up, obtain an IP address, advertise their existence and capabilities to the network, and, as appropriate, rendezvous with the sensor's controller or controllers across the home network.

This is another instance of a more generic class of automating the installation and use of services in "lightly" managed or even unmanaged networks, and intersects significantly with the objectives encompassed with SSDP and uPNP. The potential volume of such devices places this more squarely into a class of IPv6-only services, I suspect, which is a significant extension to the existing IPv4-centric uPnP frameworks.

What is needed here is a bootstrap protocol that can provide a connecting device with:

- address configuration
- routing setup
- name management and name server discovery
- discovery of other services and controllers
- security capabilities

Securing @ Home

One of the most significant issues with home networks lies in the area of security management. Host computers in a home network often want to place a very high level of implicit trust in their immediate network neighbours at the same home. Its not unusual for hosts in a home network to share printers, file servers, data, and even user profiles. Indeed, its probably commonplace. But beyond this local security domain a host should become paranoid and treat all connection attempts with suspicion. But where does the local trust domain start and stop? What is the "local" security boundary?

This is difficult to answer in an automated fashion. It's no longer the local LAN, particularly as home networks transition into routed networks. It's something to do with a local multicast scope, but that assumes that its possible to define a multicast scope that encompasses the local trust domain of the home network, and to do that we are back at the same question.

And even if you though you might have a clean answer to that question, you need to remind yourself about telecommuting. With telecommuting there is a requirement to partition out an entire local network segment and lift it out of the home environment and the home security domain and transplant it into the work security domain.

Everything @ Home

Home is certainly the new field of engagement for networked good and services. However, it's certainly one of the hardest places to play in from the perspective of attempting to deliver coherent services in a reliable and secure manner. The components are sourced from various vendors, and constructed incrementally over extended periods of times. It's an environment where legacy components need to coexist with the leading edge and the overall engineering of the environment is at best piecemeal, and perhaps more often its not engineered at all, and looks more like a random selection of technology elements assembled over an extended period of time. To make this environment work it's an environment where out-of-the-box interoperability is of paramount importance, and therefore its an environment where good standards really matter. And, perhaps unsurprisingly given these constraints, its one of the networking environments that appear to raise the most challenges. It's an unforgiving environment where there is no real substitute for simplicity and reliability in a plug and play world.

For the IETF's Homenet Working Group, there is really a lot of work to do to take a diverse set of approaches used today, add a bucketful of IPv6, and produce a coherent set of outcomes in the form of standards that support robust capable home networks that work in an unmanaged environment. By any metric that's a big ask.

Ahhh home! There really is no place quite like it!

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

Author

Geoff Huston B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. He is author of a number of Internet-related books, and was a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of Trustees of the Internet Society from 1992 until 2001.

www.potaroo.net