

What about IPv6?

Written by David Rivenburg, AD500
Tuesday, 23 February 2010 17:25 -

The intent is to migrate the mesh from IPv4 to IPv6. Here's why: 1. It's coming. In fact, it's already here. Try to go to ipv6.google.com or ipv6.netflix.com with a v4 only system and it will not be able to. Presently v6 is optional, but after v4 runs out of addresses v6 will be in your life whether you want it or not. Also, I'd rather make my mistakes now when they are not so important than later when they actually matter.

2. The large address space. Presently the 10.0.0.0/8 address space on the mesh is not large enough to allow for both unique and automatically assigned IP addresses. If it was a managed space, the available addresses would certainly be enough for the foreseeable future. But that's the point. The space is not managed. Running a DHCP server on the mesh is out of the question, and manually assigning addresses would become both an administrative nightmare and an unnecessary burden on users. Configuring a mesh node consists of setting a hostname, and soon a domain name. That is about as simple as it gets. Everything else is handled by the node itself. I'd like to keep it that way.

Now, the large addresses are going to be a pain to work with, but that's ok. Ideally you should never need to know an address. That's what DNS is for. Do any of you know the IP address of google.com? Probably not, but we all know its name and that's what we use and (network problems aside) it always works. I'd say with great likelihood that none of you have ever needed to know the IP address for google.com. And so it should be on the mesh. With node names and tactical names it should be a simple matter to navigate the mesh.

3. The large network space. With the present 10.0.0.0/8 scheme it is not possible to break down the mesh into reasonable subnets without centralized management. Again, centralized management is the antithesis of how this mesh should work. The consequence of this is that every node on the system has to know explicitly how to get to every other node by means of a routing table entry. Just as the large address space allows for automatic node address assignment, the large network space would allow for automatic subnet assignment.

If a packet coming from Austin needed to reach a node in Plano, it wouldn't have to know the specific path it needs to take to get there. All it needs to know is that all packets destined for the Plano subnet can be sent to some specific nearby node and that node is able to route the packet to its destination. That's how routing is supposed to work. Can you imagine how hard the internet would be to use if you had to know not only which website you wanted to see, but also which routers you had to use to get there? Subnets are what make intelligent, autonomous routing possible without having to enumerate every possible destination.

4. Mobile IPv6. The IPv6 protocol itself supports mobile network clients whose entry point into the network frequently changes. In fact, it competes with OLSR on this point. The routing and neighbor discovery features in IPv6 could possibly make OLSR obsolete, but at least for now OLSR will likely be easier to work with. On the other hand, mobile clients will probably see improved performance and less stuttering with IPv6 as they move from one node to the next.

5. Multicast routing - sending a single packet to multiple destinations. After all, the MM in HSMM is MultiMedia, and MultiMedia is practically synonymous with audio and video. It is much less of a burden on the network and the stream provider when multicast is used. You can do

What about IPv6?

Written by David Rivenburg, AD500
Tuesday, 23 February 2010 17:25 -

multicast in IPv4, but it is not widely implemented. It is generally not available on the internet without setting up special tunnels to make it work. In IPv6 it is mandatory that routers support multicast and it has several improvements over the IPv4 implementation.

6. I have to. My formal training is in computer engineering, but my career has been diverted into IT for the past decade. Before long an IT person without practical IPv6 experience is going to be useless. Here I am, developing an advanced networking platform on the cusp of worldwide IPv6 deployment. It would be foolish of me to not take advantage of the benefits of IPv6.

73
ad500

P.S. Bear in mind that I am talking about the **mesh**, not the local networks behind the nodes. You can manage the local networks any way you like, but the mesh as much as possible needs to be autoconfiguring and require no setup or foreknowledge of the environment in which it is expected to work.