The Standard

Volume 1 No 3 Internet Protocol version 6 Newsletter December 2009

The Standard is the journal of IPv6 Now Pty Ltd, specialists in IPv6 training, transition and innovation:
ipv6now.com.au
ipv6now.co.nz

Melbourne
7-9 Dec 2009
ipv6.org.au/summit/
WELCOME!

IPv6Now and Ai Group have constructed VIC6 to help industry make the transition to the new Internet Protocol, IPv6:
vic6.net

Seed funding for the VIC6 project was provided by MultiMedia Victoria

If IPv6 were a sport the Australian team would be in trouble

Jersey, Cuba, Vatican City – what have they got in common? They've got 50% to 100% IPv6 capability. Maybe that's due to them having just a few autonomous systems, but New Zealand has many – and a creditable 18% of them are IPv6-ready. Australia? An embarrassing 7%.

Our ISPs are a timid lot, apart from the ever-innovative Internode – but even their native IPv6 service comes with no support and no guarantees. Why on earth would any business bother going to the trouble and expense of installing IPv6 on their own networks, when they can't even get an IPv6 packet out the door? It's just not good enough to blame chickens and eggs: lack of professional IPv6 transit has become a roadblock for Australia.

Luckily, there's a way of routing around this particular damage. It's a world-class technology, fully compliant with IETF standards. Lightweight, easy, secure, reliable and able to carry tens of thousands of simultaneous IPv6 connections without missing a beat. It's an IPv6 tunnel server.

Purists may shudder at tunnels, but the sad fact is: if you're waiting for your ISP to bring you native IPv6 any time soon - don't hold your breath! The only real option in Australia at this time is tunnels. But we're not talking experimental take-it-or-leave-it services here.

We're talking *Independent IPv6* tunnels: professional, business-grade, customer support, guaranteed-service-level IPv6, ready for anywhere from home offices to enterprise networks. This edition of *The Standard* looks at tunnels in more detail. Feel free to contact me to discuss anything that interests you.

And welcome, everyone, to the Australian IPv6 Summit 2009.

Michael Biber – Professional Services Manager, IPv6Now

Contents

What sort of IPv6 do you want?

Independent IPv6

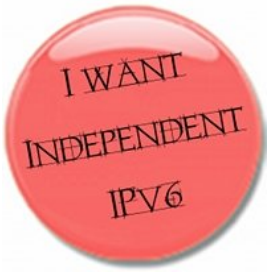

Your network, email and hosting services are supplied by an Internet Service Provider. You want to access IPv6, and your ISP or equipment vendor isn't doing it any time soon. What do you do? Route around that logjam! Get Independent IPv6 access immediately, with no changes required to equipment or ISP services. Business-grade Independent IPv6 is suitable for Windows, Mac or Linux systems, scales easily to multiple users and multiple sites, and is protected by Service Level Agreements.

Connected IPv6

You're a communicator: being connecting to the world is what matters. Your websites and mail services are professionally hosted, and now you want to reach everyone, not just the IPv4 people. What can you use? Get a Tandem Service, allowing IPv4 and IPv6 to work together cooperatively. Try EasyWeb6, which makes your content visible to both IPv4 and IPv6 users with no equipment or ISP changes, or 6Now Hosting & Email on commercial dual-stack systems, all protected by Service Level Agreements.

Integrated IPv6

You or your staff run your own networks and services. You're clued-up about IPv4 but wondering whether your networks, databases or webservers can handle IPv6. How do you integrate IPv6 into existing systems? Use Verification Testing – it provides network auditing, vendor-neutral advice, IPv6 installations, and customised testing with the only publically-available IPv6 testbed in Australia. Don't lose customers because you can't talk to the growing IPv6 networks in China and the developing economies.

Power IPv6

Knowledge is the source of all power: if you want IPv6 power then you must understand the issues. IPv6Now has the training courses suitable every level of your business – upper management, company support, technical implementors. We can provide professional advice based upon decades of Internet experience, plus the resources you need to discover your many IPv6 options. Know what's coming down the track before it arrives, and better still – find out before your competitors do.

All the IPv6 you could want – ipv6now.com.au

Major Tunnel Techniques

Michael Biber, Professional Services Manager, IPv6Now

The mechanisms for deploying IPv6 across networks range from a complete 'big-bang' conversion, through dual-stack approaches, or running IPv6 within existing IPv4 infrastructure. This article looks into the latter approach: the options that exist and some comments on the pros and cons.

Perhaps the most desirable outcome would be to have a flag day where all IPv4 is turned off and IPv6 is turned on in everyone's network simultaneously, but this is completely impracticable for the whole world of Internet users, and also largely impractical for any individual user.

As Geoff Huston has said "The most common scenario of IPv6 enterprise and end user deployment is still one of isolated islands of IPv6 in an ocean of IPv4. The relative isolation of these IPv6 networks means that the only practical way of supporting a connected IPv6 Internet is to treat the IPv4 network as a lower layer connectivity tool for IPv6 and tunnelling IPv6 across the IPv4 'undernet'". [Ref 1].

Most ISPs and supporters of IPv6 recommend a dual-stack approach, where native IPv4 and native IPv6 capabilities are built into the equipment and networked applications. Where IPv6 is available, this is used, and where IPv4 only is available, this is used.

An example is a web browser (such as Firefox) where IPv6 DNS lookups are made first, and if there's no response from a AAAA record in the name server, this is followed by the lookup of the IPv4 A record. This approach seems to work well enough and the additional overhead does not appear to be too onerous.

There are however, questions as to how well this will scale across the whole spectrum of Internet activity. Generally, IPv6 transition mechanisms are part of one of two classes: translation or tunnelling. Translation involves transforming an IPv6 packet into an IPv4 packet and vice-versa, but in general it is not recommended as a transition mechanism. Tunnelling is the act of putting an IPv6 packet within an IPv4 packet so it can transit the IPv4 parts of the network.

Tunnelling (IP in IP) can be achieved in many different ways, but at a fundamental level this involves wrapping the IPv6 packet in an IPv4 envelope (frame), so that the IPv4 Internet can route the whole datagram to its destination. The reverse can also occur. where IPv4 data packets can be wrapped in IPv6 headers and routed through an IPv6-compliant network.

Tunnelling is a versatile tool and there are a number of techniques being used for this purpose. As discussed in Ref 1, a tunnel can be configured as a virtual interface on an individual host system, allowing individual hosts to configure a hop over the local network to make a virtual direct connection to a remote access point.

A tunnel may also be configured as a virtual interface on a router, allowing a network to use this virtual connection to directly connect to another network, or act as a remote access concentrator for hosts or end sites. Among the most common tunnelling protocols are 6to4, Teredo, ISATAP and TSP.

6to4

The 6to4 approach is described in RFC3056, RFC 3068 and RFC3964. 6to4 is a framework of automated point-to-multipoint tunnels that addresses the same general scenario as the manually configured tunnel environment, namely a collection of "islands" of IPv6 that are interconnected by the public IPv4 network.

The change with 6to4 is that each site is no longer configured with a single tunnel to some tunnel hub, but is in effect configured with the local end of a point-to-multipoint tunnel, so that each 6to4 site is capable of sending and receiving packets from the entire set of 6to4 end sites, using the same local tunnel endpoint.

In effect all 6to4 sites form a fully meshed overlay network, with automatically managed 6to4 tunnels to every other 6to4 site. Note that standard 6to4 requires the underlying IPv4 network to support IP multicast, or increasingly, IP anycast.

Teredo

Like 6to4, Teredo (described in RFC 4380) is both an address assignment and tunnelling technology. Teredo is commonly used to work through NATs and uses a version of the STUN active probing approach to determine the type of NAT.

Teredo uses concepts of "clients", "servers" and "relays". A *Teredo client* is a dual stack host that is located in the IPv4 world, possibly behind a NAT. A *Teredo server* is an address and reachability broker that is located in the public IPv4 Internet, and a *Teredo relay* is a Teredo tunnel endpoint that connects Teredo clients to the IPv6 network. The tunnelling protocol used by Teredo is not the simple IPv6-in-IPv4 protocol 41 used by 6to4.

NATs are sensitive to the transport protocol and generally pass only TCP and UDP transport protocols. In Teredo's case the tunnelling is UDP, so all IPv6 Teredo packets are composed of an IPv4 packet header, a UDP transport header, followed by the IPv6 packet as the tunnel payload.

Note that Teredo relies on ICMP to complete an initial protocol exchange and confirm that the appropriate NAT bindings have been set up. Firewalls that routinely discard ICMP messages will disrupt communications with Teredo clients.

ISATAP

ISATAP (Intra-Site Automatic Tunnel Addressing Protocol) described in RFC 5214, is an IPv6 transition mechanism meant to transmit IPv6 packets between dual-stack nodes on top of an IPv4 network. Compared with 6to4 use of IPv4 multicast, ISATAP uses IPv4 as a virtual nonbroadcast multiple-access network (NBMA) data link layer, so that it does not require the underlying IPv4 network infrastructure to support multicast.

ISATAP utilizes an innovative principle in which an IPv4 network emulates a logical IPv6 subnet to a set of ISATAP hosts. This principle allows all ISATAP nodes, no matter where they are located on the IPv4 network, to automatically tunnel to each other for IPv6 connectivity and reach other IPv6-capable networks or the IPv6 Internet through an ISATAP router. [Ref 3].

A new initiative is underway to create a next-generation version of ISATAP that support IPv6 autoconfiguration, adaptation to diverse tunnel MTUs and IPv6 multicast. This next-generation version is named Virtual EThernet (VET) and will use the Subnetwork Adaptation and Encapsulation Layer (SEAL).

TSP

So far, all of these tunnelling mechanisms all have one fundamental problem; a host is not guaranteed to have a fixed IPv6 address. This makes it difficult to establish end-to-end trust or to act as a server and run any type of publicly advertised services. These protocols, especially 6to4 because it uses anycast, are harder to debug as the return path of the packets may be different than their forward path.

However, Tunnel Set-up Protocol (TSP) based tunnels, provided through redundant tunnel broker platforms (e.g. gogoServer), offer session authentication capabilities, improving their security. They also facilitate the creation of accounting records and detailed network management monitoring for each user. The IPv6 addresses and prefix assigned by these mechanisms are fixed, making them suitable to run trusted, public network services.

TSP is a lightweight protocol, designed to be suitable for all types of hosts, including small embedded devices. It is invoked though a small client application loaded on individual hosts or various network devices. As one issue with negotiated tunnelling methods is that all the IPv6 traffic goes through a single server, but this is not important with redundant services.

In addition, clients have the opportunity to implement improved security though the control and inspection of all IPv6 traffic flows, into and out of their networks, through well defined portals.

TSP is the protocol underlying gogoServer, used by IPv6Now for its business-grade Independent IPv6 tunnels with SLA guarantees.

References

Ref 1 Geoff Huston, February 2008, The ISP Column, 'IPv6 Transition Tools and Tui'
Ref 2 http://www.isatap.org/
Ref 3 Microsoft/Cisco White paper 'IPv6 Integration Paper Series: Manageable Transition to IPv6 using ISATAP, An Executive Overview', May 2006
Ref 4 www.gogo6.com/4105/description.asp?product_id=172
Ref 5 draft-blanchet-v6ops-tunnelbroker-tsp-04

The IPv6 Tunnel Zoo

Kate Lance, Communications Manager, IPv6Now

	6in4 static	6in4 heartbeat	6to4	6rd	ISATAP	Teredo	AYIYA	L2TP	TSP/ DSTM	TSP
Protocol basis	41	41	41	41	udp	udp	udp	udp	udp	udp
Able to traverseNATs	N	N	N	N	N	Y	Y	Y	Y	Y
Address space	any	any	2002::/16	any	any	2001::/32	any	any	any	any
Does IPv6 over IPv4	Y	Y	Y	Y	Y	Y	Y	Y	N	Y
Does IPv4 over IPv6	N	N	N	N	N	N	Y	Y	Y	Y
Has dynamic endpoint	N	Y	Y	Y	Y	Y	Y	Y	Y	Y
Packet overhead	low	low	low	low	med	high	high	high	high	med
Requires multicast	opt	opt	N	opt	N	N	opt	opt	opt	opt
Provides authentication	N	Y	N	N	N	N	Y	Y	Y	Y
Prefix delegation	opt	opt	48	opt	N	N	opt	opt	opt	opt
Allows traffic accounting	opt	opt	N	opt	opt	N	opt	opt	opt	opt
Reverse DNS	opt	opt	opt	opt	opt	N	opt	opt	opt	opt
RFC	4213		3056		5214	4380		2661		

Table data from SixXS - http://www.sixxs.net/faq/connectivity/?faq=comparison

No Nat Traversal

6in4 (static) uses tunneling to encapsulate IPv6 traffic over explicitly-configured IPv4 tunnels with fixed endpoints.

6in4 (heartbeat) provides 6in4 with dynamic tunnel endpoints using the heartbeat protocol, which signals the other side of the tunnel with its current endpoint location. Has authentication.

6to4 is router-to-router automatic tunnelling that uses prefix 2002::/16. An isolated IPv6 site assigns itself a prefix of 2002:V4ADDR::/48, where V4ADDR is the IPv4 address on the egress router, then IPv6 traffic is encapsulated as usual. The endpoint must have a public IPv4 address, and the network has to support multicast.

6rd (rapid deployment) is like 6to4, except that a 6rd service provider uses its own IPv6 prefix rather than the fixed 6to4 prefix 2002::/16.

ISATAP (Intra-Site Automatic Tunnel Addressing Protocol) connects dual-stack IPv6/IPv4 nodes over IPv4 networks *within* trusted sites, using udp. Unlike 6to4 it does not require multicast.

NAT Traversal

Teredo allows IP6 connectivity to nodes behind IPv4 NAT devices, by encapsulating IPv6 packets within IPv4 udp datagrams, which can then be routed through the NATs. Hence IPv6 hosts behind NATs can be Teredo tunnel endpoints without a public IPv4 address.

NAT Traversal + Authentication

AYIYA (Anything In Anything) is a tunneling protocol that can be encapsulated within any other protocol. It uses authentication tokens, allowing multiple identities on the same tunnel endpoint. The endpoint can be changed on the fly, so AYIYA can be applied in multihoming and mobility solutions.

L2TP (Layer Two Tunneling Protocol) is used in virtual private networks in conjunction with IPsec security.

TSP/DSTM (Dual Stack Transmission Mechanism) is for tunneling legacy IPv4 host traffic over predominantly IPv6 networks. (Obviouslynot a lot of call for it yet.)

TSP (Tunnel Setup Protocol) is a signaling protocol to set up tunnel parameters between two tunnel end-points. TSP is implemented as a tiny client code in the requesting tunnel endpoint. The other endpoint is the server that provides the tunnel service. Parameters always negotiated are:

- Authentication of users
- Encapsulation - IPv6 over IPv4 , IPv4 over IPv6, IPv6 over UDP-IPv4 for NAT traversal
- IP address assignment for tunnel endpoints
- DNS registration of IP endpoint address

Other parameters that may be negotiated:

- Tunnel keep-alive
- IPv6 prefix assignment when client is a router
- DNS delegation of the inverse tree, based on the IPv6 prefix assigned
- Routing protocols

Zooming in on IPv6 Implementation

Kate Lance, Communications Manager, IPv6Now

The closer you zoom in on any major deployment the more fractal-like it becomes: details within details. At the highest level, deploying IPv6 is something like this: [A] Get IPv6 connectivity to the ouside world, and [B] Implement IPv6 on internal systems.

[A] is not as hard as it might seem. Ask yourself this:

1. Does your ISP have IPv6 with customer support and guaranteed service levels?
   If you're in Australia, that's a NO. (Enjoy the wait.)
2. Do you want IPv6 connectivity without having to wait, including support and SLAs?
   If that's a YES, get business-grade Independent IPv6 from IPv6Now.

There – external IPv6 done! (Sorry about the product placement, but sadly it's true.)

[B] is more complex and long-term. At this level you have to:

1. Decide upon project parameters and management.
2. Educate yourself and company personnel about IPv6.
3. Discover the IPv6 status of software and hardware.
4. Use qualified personnel for a phased implementation.

[B] expands into more detail:

1. **Project parameters and project management**: this will depend upon specific business size and complexity, however IPv6 integration must be managed by a qualified person, as it is too important to get wrong.

2. **IPv6 education**: the are many levels to IPv6, it is *not* simply a technical matter. Keep in mind that last-minute training will lead to expensive competition for training resources:

   a) Management IPv6: needed in respect of strategic, geopolitical and customer aspects.
   b) Business IPv6: must be understood by general personnel, as it will affect customer interfaces, business applications, internal and external support, call centres.
   c) Technical IPv6: training for technical implementors is fundamental.

3. **Software, hardware and IPv6 compatibility**:

   a) Audit software and hardware to find out IPv6 status.
   b) Determine requirements for upgrading or replacing.
   c) Try to purchase during the natural upgrade cycle. Don't leave to the last minute as that will mean the least cost-effective purchases.
   d) Rewrite any in-house applications that use IP addresses for full IPv6 compatibility.

4. **Qualified personnel and phased implementations**:

   a) Use trainers to bring staff up to speed on IPv6.
   b) Use trained in-house or external personnel for IPv6 software and network deployment.
   c) Commence IPv6 implementation.

Now item 4 expands into more detail. Areas affected by IPv6 to focus upon:

1. Address deployment plan
2. Routing protocol set-up
3. Setting up IPv6 transit and peering
4. Determination of IPv6 routing policy
5. Traffic engineering options
6. Load balancing for service platforms
7. Enabling existing platforms for dual stack
8. Firewalls and security framework
9. Maintain legacy facilities and applications
10. Domain Name System
11. Email services
12. Web services
13. Application support
14. Customer-visible services
15. Customer premises equipment
16. Billing and accounting
17. Disaster recovery and business continuity
18. Operations and management tools

Each of these areas will have its own levels of complexity, based upon your own business structure and staff knowledge. The time and effort taken will depend upon those factors, but early adopters have already reported that IPv6 transition is not as daunting as it might first appear. Also, as more vendors implement IPv6 over time, equipment and software aspects will become easier, although training will always be essential.

Like *any* project, IPv6 may go smoothly - or have a few bumpy moments. But once it's complete, sit back and enjoy the fruits of your efforts: a flexible, clean, and easy-to-manage business foundation, with a potential for innovation that was not possible with IPv4.

KANZ09 Confirms a 2,100 Year Old Principle

Kevin Karp, Business Manager, IPv6Now

My attendance at two recent events, each discussing advanced technology computing, clarified for me an age-old principle: high technology is there to serve the very human needs of ordinary folk. It didn't seem to matter that the two advanced computing technologies were separated by 2100 years of history.

I went to a talk by Professor Robert Hannah on the Antikythera Mechanism (http://en.wikipedia.org/wiki/Antikythera_mechanism) – the oldest known complex scientific calculator, if not the first analog computer. Prof Hannah discussed research on the device, including the replica made by my old University of Sydney Computer Science lecturer, the late Allan Bromley.

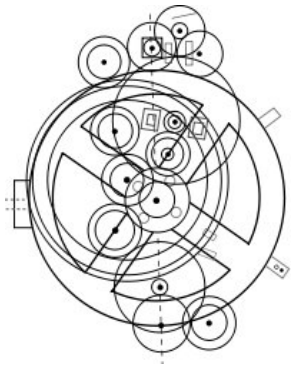

He concluded by speculating on why the mechanism was built by the ancients, downplaying the usual explanations of an astronomical instrument, a civic calendar, a calendar to track the next Olympic games or even a weather prediction device. He felt its real purpose was for everyday astrological predictions, just like the popular columns in today's magazines.

Then at the KANZ09 Broadband Forum in early November I learnt of the huge advances being made in Korea with the deployment of their broadband network: bandwidth speeds of 300Mb download and 100Mb upload now, to be doubled in 2010, to be rapidly replaced with Gb speeds in 2011. And one of the major applications today for this enormous capacity: *interactive residential karaoke*!

The message for me in considering the uses for the new IPv6-enabled broadband network for Australia is *don't forget the interests of ordinary people* – after all, there's 2,100 years of precedent ...

Victorian Industry Collaborative IPv6 TestNet

The purpose of the VIC6 project is to provide an industry facility in the form of a distributed IPv6 testbed that will enable Victorian organisations to examine IPv6 implementation, integration and innovation, with a view to deploying the protocol within their networks, products and services. The project has four phases:

- Phase 1: VIC6 Construction
- Phase 2: VIC6 and the Victorian Health Sector
- Phase 3: VIC6 and the Victorian Financial Services Sector
- Phase 4: VIC6 and the Victorian Automotive Manufacturing Sector

In just 12 months the VIC6 project has established a new TestNet for industry to develop its capabilities with the new Internet technology, Internet Protocol version 6 (IPv6). The TestNet environment was based upon on IPv6Now's extensive experience with Interne technologies. It provides industry with a reliable but flexible base to explore IPv6, in areas ranging from skills development, through transition strategies and staging solutions, to product and process innovation.

Industry outreach during the 12 months has included general industry communication and consultation, as well as targeted involvement of working groups in health, financial services and automotive manufacturing. As a result of this engagement, VIC6 has been able to analyse the

business needs of these sectors and generate key IPv6 demonstrations that meet those needs.

Key Developments

Phase 1: Construction
The core node of VIC6 at Ballarat continues in operation. Agreement to host a second core node at Bundoora has been completed with La Trobe University, and we anticipate installation of this node in the near future. Discussions have begun with the Institute for a Broadband-Enabled Society (IBES) to host a third core node and we anticipate further positive discussions in the near future. Planning for the VIC6 network operations centre (NOC) has commenced, with hardware currently scheduled to be deployed.

Phase2: Health
VIC6 has demonstrated key operation of a general purpose network implementation that can support data, voice and video using IPv6 for health applications. A videoconference demonstration has been successfully completed between Ballarat and the Centre for Eye Research in Melbourne.

Simpler remote access over IPv6 has been demonstrated with Eyescan, a new eye health service delivery business. The Centre for Health Innovation has been established as the first VIC6 Industry Node after enabling IPv6 on its network, as a future foundation for innovate health solutions. Further work will continue, for instance, with the Australian Health Messaging Laboratory to demonstrate operation of the HL7 health messaging standard over IPv6.

Phase 3: Financial Services
Material for the business analysis of the needs of the financial services sector was obtained by VIC6, despite the sector's considerable distraction regarding the global financial situation, and additional industry outreach and activities are continuing. As a result of business discussions, VIC6 has defined the potential for a range of demonstrations related to the needs of the financial services sector. These include:

- The relative performance under IPv4 and IPv6 of IPSec (the standard framework for securing Internet communications by authenticating and encrypting each IP packet of a data stream)
- An IPv4 banking system using IPv6
- An access authentication model over IPv6
- Demonstrating eVision's MXC B2B platform technology over IPv6 (MXC is the first all-in-one, on-demand, Internet service for hosting business exchange of electronic documents and messages).

Phase 4: Automotive Manufacturing
VIC6 is consulting with the Victorian automotive manufacturing sector on its business needs in relation to IPv6. The week of 16-20 November 2009 has seen the State of Victoria become the national focus of activity in Intelligent Transportation Systems. VIC6 has participated in these events, and also held its own designated industry consultation session on 18 November.

Briefing and discussions have already been held with the Federation of Automotive Parts Manufacturers (FAPM). VIC6 has engaged with Embedded Systems Australia, and the AusDSRC Industry Cluster (a VIC6-based demonstration is designated as potential project 6).

Future Activities: RFID Demonstration
Following strategic discussions with Ai Group and the RFID Working Group, VIC6 is planned to be designated as the venue for continuing industry outreach and practical demonstrations. A new area of the VIC6 wiki will be designated for RFID discussions, and businesses with an interest in RFID practical demonstrations will have access to the facilities of VIC6, particularly to demonstrate operation of RFID technologies using IPv6 as a platform for data communications.

Further Information
Information on VIC6 activities is available from the VIC6 website at: www.vic6.net. Contact: Tony Hill, Managing Director, IPv6Now, tony@ipv6now.com.au, 02 6161 6607.

VIC6 TestNet Facilities

Karl Auer, Technical Manager, IPv6Now

How can your organisation test for itself whether or not its facilities are going to be compatible with IPv6? You could build your own IPv6 laboratory with techies and expensive equipment, at a significant cost in both time and investment. Or perhaps you could choose to use a state-of-the-art IPv6 testbed facility, courtesy of the Victorian Government.

Due to the foresight of MultiMedia Victoria, we

now have the VIC6 TestNet - a fully functioning IPv6 network for testing IPv6 compatibility of hardware and software, *without* any risk of damaging your vital production systems. So just what has VIC6 got to offer?

The Basics

The VIC6 TestNet has native IPv6 connectivity to the IPv6 Internet. Your testing is not limited to a closed, local network; you can test with other organisations or individuals around the state, the country or the world.

You have a fully-fledged IPv6 and IPv4 dual-protocol network. A Cisco router connects the network to the IPv6 Internet, and a Cisco switch provides 10/100/1000 connectivity within the network. VLAN (Virtual Local Area Network) technology allows the 48 ports on the switch to be grouped into various different subnets with different properties. If your testing requires it, we can provide a customised VLAN (or more than one).

All of the above has a place to live, and expert support personnel to tend it. Visitors to the node can take a seat in a comfortable workspace and connect their laptops or other equipment to the VIC6 node directly. For more complex testing or experimentation, the node has rack space in a professionally maintained data centre - so if you want to test a router or a switch, or you have a server you'd like to install for a while, we can rack it and stack it for you. Lateral Plains also provides the virtualisation expertise underpinning services like DHCP and DNS.

Each node has one or more servers providing DNS, DHCP, SSH (secure shell), etc, to VIC6 users. VIC6 partners get remote IPv6 access to the VIC6 TestNet via IPv6Now's commercial tunnel broker in Australia, the only one in the world with guaranteed service level agreements.

Each VIC6 node has industrial strength DNS, DHCPv4 and DHCPv6 services. Instances of the ANS, Vantio and DCS3 products handle the normal DNS and DHCP requirements of each node, plus each node has an instance of ANS and an instance of DCS3 available for you to experiment with.

VIC6 has sophisticated testing and simulation equipment for the Ballarat node. The Spirent Test Centre is expected to be available by arrangement to any VIC6 partner wishing to perform traffic simulation, test a network design and so on, with basic familiarisation assistance to VIC6 members free of charge. More comprehensive assistance (and sales of hardware) can be provided to VIC6 members at a generous discount.

The whole VIC6 network is watched over 24/7. PPS provides NOC (network operations centre) services to monitor the network, back up critical data and provide an escalating response in case of problems.

VIC6 VLANs

The VIC6 TestNet has **four standard Virtual Local Area Networks** available for use, plus one management VLAN and one services VLAN. The DNS, DHCP and other similar services are in the services VLAN. The four "user" VLANS are:

- **Public**: with full access to *and from* the Internet
- **Protected**: with full access *to*, but limited access *from* the Internet
- **No Internet**: with no access to or from the Internet
- **Isolated**: with no access at all to anything outside the VLAN

Devices can be connected to the appropriate VLAN as desired, and **custom VLANs** with particular characteristics can be arranged.

In the public, protected and isolated VLANs, the router advertises a /64 IPv6 subnet, so attached equipment can **use autoconfiguration** to acquire a globally routable IPv6 address.

In the public, protected and isolated VLANs, DHCPv4 and DHCPv6 are available. Attached equipment can acquire **globally routable addresses via DHCP** if so configured. If statically allocated addresses are desired, these are supported too, and the DHCP servers support **stateless IPv6 autoconfiguration**.

Each VIC6 node has an additional instance of the Nominum DCS3 **DHCP server available for experimentation**. In conjunction with a custom VLAN, VIC6 members can obtain total control over the address distribution mechanisms.

Each VIC6 node has its own DNS, secondaried on IPv6Now nameservers. All services and servers are addressable by name or by IP address; most are dual-stack.

Each VIC6 node has an additional instance of the Nominum ANS **nameserver available for experimentation**; DNS views can be created for specific purposes, using whatever nameserver information you desire (although the exchange of DNS information between these servers and the global Internet can only be permitted if consistent with usual DNS practices and norms).

An **ssh server is provided**, and each VIC6 partner may have one or more accounts on this server. From this ssh server, connections to other equipment within the VIC6 node can be made. This provides a simple springboard, protecting machines that should not be directly accessible from outside the VIC6 TestNet. Each VIC6 node can also provide, by arrangement, a **virtual server** to VIC6 partners that need to run software within the VIC6 node.

Would you like to know more? Sign up for involvement with the VIC6 project at **vic6.net** and find out how you can use the VIC6 TestNet to better understand your own organisation's IPv6 issues and options, and the most cost-effective implementation pathways to adopt.

VIC6 was built with the assistance of sponsors Neighbourhood Cable, Cisco Systems, Lateral Plains, IPv6Now, Nominum, Matrium, and PPS Internet.

Social Networking Plus IPv6 Equals?

Kevin Karp, Business Manager, IPv6Now

Social networking is a highly popular Internet functionality permitting spontaneous, dynamic interaction between individuals and groups. It is best known for its friendship communities, but it also has potential as a tool for education, government service delivery, business links and medical communication. Unfortunately, social networking also shows some negative aspects, such as cyberbullying, identity theft, lack of auditability, lack of traceability and infringements of privacy.

IPv6 is a new Internet protocol permitting greatly expanded capabilities to support peer-to-peer interaction at both a technical and personal level. It provides features such as on-the-fly dynamic allocation of personal DNS entries and individually allocated IP addresses, which provide improved peer-to-peer activity, greater portability of identity, and simpler and more reliable identity tracking.

What do you get when you bring social networking and IPv6 together?

First, IPv6 can help alleviate some social networking negatives by:

- Improving logging accuracy
- Providing better identity validation
- Simplifying auditing and tracing
- Making illegal activity harder to hide

Second, IPv6 can technically enhance social networking by:

- Increasing the flexibility of peer-to-peer interactions
- Dramatically expanding the capacity of social networks
- Improving performance by allowing more open simultaneous ports
- Providing more functiionality for social network technical developers

Adding another ingredient: the introduction of SAML (Security Assertion Markup Language) permits identification and authorisation to be performed across internet web sites and applications. This cross-site authorisation provides for a Single Sign-On (SSO) experience. SAML can enhance the integration of social networking sites.

The abundance of IPv6 allows for the permanent assignment of IP addresses, and in combination with Single Sign-On and positive personalised DNS entries, creates third cluster of benefits:

- convenient site interaction for the user
- a single point of userid/password admin
- more auditable activity for administrators

One final overlap between IPv6 and social networks! gogo6 (the company previously known as Hexago) has started gogoNET, a social network and services for professionals to go v6. IPv6Now is proud to be one of the four worldwide PoPs that support the network. See http://gogonet.gogo6.com/

IPv6 Understanding – We all Need It

Michael Biber, Professional Services Manager, IPv6Now

For those of you looking at Windows 7 be aware that it installs with automatically configured Teredo tunnels. Your PC will be connected to the IPv6 Internet automatically: *do you know what this means for you or your business?*

IPv6 is not just a technical development, it is a foundation of business strategy for many years to come. Any organisation that uses computer networks can expand, innovate, develop new markets and gain a competitive advantage with IPv6.

IPv6Now training currently take place in Sydney, Canberra and Melbourne, and may be tailored to other requirements or locations. Courses are independent of vendor bias.

Contact Michael Biber for details: 0412 058 808 or michael@ipv6now.com.au

Everyone needs to know about IPv6 at a level that suits their responsibilities. IPv6Now offers three types of course, for decision-makers, business support and technical implementors:

Strategic IPv6 for Decision-Makers

Strategic direction must begin at the top. Executive Directors, Heads of Information Technology, Finance, Sales, Development, Logistics and other essential departments need to understand the global impact of IPv6, the political state-of-play, and the commercial implications of IPv6 as a whole, before selecting long-term strategies and investments.

Strategic IPv6 Course (2 hours)

- What is IPv6?
- Current Situation
- Advantages of IPv6
- Impact of IPv6
- IPv6 in Australia
- IPv6 Business Strategies

Business IPv6 for Company Support

Companies must support products and processes for customers and for their own personnel – and even those that don't use IPv6 will have to interact with customers, partners and competitors that do. People who will benefit from this level of IPv6 knowledge are call centre responders, sales staff, website developers, applications programmers, technical writers, computer administrators and database maintainers.

Business IPv6 Course (1 day)

- What is IPv6?
- Problems with IPv4
- New Features of IPv6
- IPv6 Address Format
- IPv6 Packet Overview
- IPv6 Security and Mobility
- Programming with IPv6
- Transitioning to IPv6

Technical IPv6 for Implementors

At this level no-one cannot afford to get IPv6 wrong. The course provides full details of IPv6 formats, techniques and interactions, including practical hands-on exercises in software and hardware. It is ideal for training IT managers, computer security staff, network designers, programmers, and trouble-shooters, and system and network implementors and administrators. The course refers to Windows, MacOS, Linux and Cisco IOS.

Technical IPv6 Course (3 days)

- IPv6 vs IPv4
- IPv6 Addresses
- IPv6 Packet Format
- IPv6 Security
- Mobile IPv6
- DNS and IPv6
- IPv6 Autoconfiguration
- Internetworking IPv6
- Dynamic Routing Protocols
- Network Management in IPv6
- Transition Mechanisms

www.ipv6now.com.au/courses.php

IPv6 Summit Games Page
Play IPv4 Bingo!

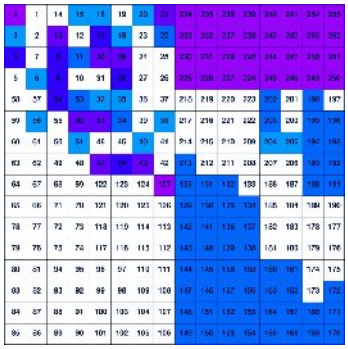

IPv4 address space, end 1994

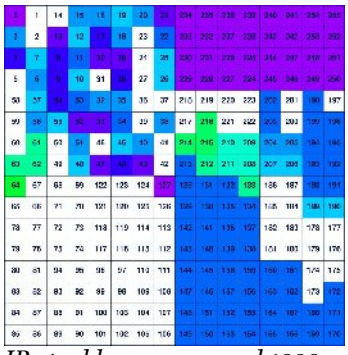

IPv4 address space, end 1999

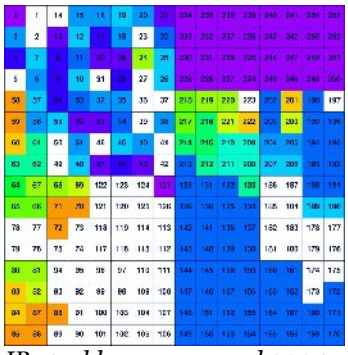

IPv4 address space, end 2004

Use a red pen to fill in the blocks allocated over the next two years – watch IPv4 space dwindle and disappear before your very eyes ... which /8 will be the next to go?

0	1	14	15	16	19	20	21	234	235	236	239	240	241	254	255
3	2	13	12	17	18	23	22	233	232	237	238	243	242	253	252
4	7	8	11	30	29	24	25	230	231	226	225	244	247	248	251
5	6	9	10	31	28	27	26	229	228	227	224	245	246	249	250
58	57	54	53	32	35	36	37	218	219	220	223	202	201	198	197
59	56	55	52	33	34	39	38	217	216	221	222	203	200	199	196
60	61	50	51	46	45	40	41	214	215	210	209	204	205	194	195
63	62	49	48	47	44	43	42	213	212	211	208	207	206	193	192
64	67	68	69	122	123	124	127	128	131	132	133	186	187	188	191
65	66	71	70	121	120	125	126	129	130	135	134	185	184	189	190
78	77	72	73	118	119	114	113	142	141	136	137	182	183	178	177
79	76	75	74	117	116	115	112	143	140	139	138	181	180	179	176
80	81	94	95	96	97	110	111	144	145	158	159	160	161	174	175
83	82	93	92	99	98	109	108	147	146	157	156	163	162	173	172
84	87	88	91	100	103	104	107	148	151	152	155	164	167	168	171
85	86	89	90	101	102	105	106	149	150	153	154	165	166	169	170

Since we first published this plot in August 2009, blocks 2 and 46 have been allocated. Only 26 remain, and the last of them is expected to go in 2011.

© IPv6Now Pty Ltd, ACN 126 460 955
Editor: Dr Kate Lance (editor@ipv6now.com.au)
See earlier editions of *The Standard* at *ipv6now.com.au/standard.php*
ipv6now.com.au ipv6now.co.nz