IPv6 for e-Business - http://www.ipv6.org.au/

Activity 4.1 Underpinning Infrastructure

See http://www.ipv6.org.au/underpin.html

This is a summary of major issues relating to IPv6 in the areas of address allocation, DNS operations, transition support, e-commerce and security keys. (See the glossary of terms below, with links to further information.)

Introduction

IPv6 the protocol is well-defined. Less well defined are the critical support infrastructures. Does the process for actually getting IPv6 addresses from registrars work? Can DNS entries that map IPv6 addresses to useful names be made and relied upon? What support is out there for organisations transitioning to IPv6? Can existing e-commerce infrastructures handle IPv6? What problems are posed by the new security functionality in IPv6? This document identifies the major infrastructure issues in each of these areas and describes their current status and likely resolution.

Part 1. Address Allocation

1.1 Discussion

For IP addresses to be as useful with IPv6 as with IPv4, they must be globally unique [1]. The top level of control rests with the Internet Assigned Numbers Authority, IANA. The task of actually distributing addresses rests with a number of Regional Internet Registries or RIRs.

Making address space available for further distribution is known as **allocating** it. Making address space available for use in an operational network is known as **assigning** it.

Each RIR is responsible for a geographical area - Europe, Asia-Pacific and so on. Each RIR is allocated a (very large) block of addresses by IANA; these are then either assigned in smaller blocks directly to organisations that need them or allocated to National Internet Registries (NIR) or Local Internet Registries (LIR) for further distribution. This is a very simplified description of the arrangement; for more information see [2].

IPv6 addresses are essentially free, or at least extremely cheap [3]. This is a positive thing, in that the cost of addresses is not a barrier to the uptake of IPv6. On the other hand, there is no incentive for the users of IPv6 address space to be parsimonious; this could lead to 'land grabs' of inappropriately large amounts of IPv6 address space - whether through ignorance, greed or carelessness. While the IPv6 address space is vast, simply assigning it in vastly large chunks could still exhaust it. This is clearly something to be avoided, and the RIRs have policies in place that are designed to minimise this problem.

Once an organisation has been assigned addresses, those addresses need to be made 'visible' within the Internet. This is done by publishing a route to the address block being used. The goal is to publish routes only for large blocks, to minimise the amount of routing information that must be managed. Thus, an organisation with a large address block publishes one route for the entire block, and manages the routing within that block itself. This is known as route aggregation, and address space that can be aggregated in this way is known as **provider aggregatable** (PA) address space.

Some organisations want to have more than one 'upstream' connection. Typically this is done to provide redundancy. That is, if one connection fails, communication can continue over the other connection. This is known as **multihoming**. However, for any given PA address block, there is only one route to that block. If an organisation has addresses from two sources, it can publish two routes, but any given address remains accessible only via one route. PA addresses thus cannot provide the desired redundancy.

The standard solution is to have so-called **provider independent** (PI) or 'portable' address blocks. These are addresses that are not tied to any one access provider. The downside to this solution is that the routes to such spaces must be stored in many places - there is no longer just one route to that address block. If thousands or even millions of organisations use such address blocks, the sheer volume of routing data that needs to be stored around the world can become a very serious problem.

Technically, the standard solution can be used with IPv6, as it has been used for many years with IPv4. However, due to the routing issues, IPv6 address assignment policies forbid the allocation of portable address space. These policies are now being relaxed to allow IPv6 multihoming [4], while a technical fix to the routing issue is being sought.

IPv6 address space has several advantages over IPv4 address space, quite apart from the much larger number of addresses available. One such advantage is the ability to renumber a network relatively easily. With IPv6, interfaces can have several addresses, and a new address can be added and used without invalidating the old one. This allows an almost interruption-free transition from one address plan to another. Changing from one provider to another typically involves changing only the prefix part of the addresses in use; subnetting and host addressing can remain unchanged.

1.2 Summary

The situation with address allocation is generally good. People can certainly get enough addresses. There is a small issue with requests for too many (or, more rarely, too few) addresses, but the renumbering functionality in IPv6 makes fixing such errors relatively painless.

A major issue is multihoming. Organisations (typically carriers and Internet service providers, but also larger corporates) who want to do multihoming do not care about the routing problems, they simply want to be able to multihome, and may not move to IPv6 if they cannot do so. The current policy proposal to issue portable addresses to organisations which are, or plan to be, multihomed will do a very great deal to remove this issue.

1.3 Next Steps

Anyone wishing to implement IPv6 can get IPv6 address space right now and proceed.

Organisations wishing to multihome can still do a lot to get started, but will have to wait until the proposed policy on PI address space is implemented. Fortunately that is due to happen in 2007 for the Australia-Pacific region. Alternatively, such organisations can start now with PA address space and plan to renumber when portable address space becomes available.

Part 2. DNS Operations

2.1 Discussion

The Domain Name Service (DNS) is the mechanism that translates addresses into human-readable, useful names. IPv6 offers a comprehensive DNS security model, and, like IPv4, can be secured within the DNS using DNSSEC.

IPv6 addresses are physically larger than IPv4 addresses. IPv6 addresses are 128 bits long; IPv4 addresses are only 32 bits long. Where IPv4 addresses can be written down in no more than 15 characters and are written in decimal, an IPv6 address is written in hexadecimal notation and may require up to 47 characters [5].

IPv6 addresses also have more internal structure (prefix, subnet and host parts, for example), and many addresses have special meanings. Taken together, these things mean that for all practical purposes, IPv6 addresses are no longer humanly manageable except in the smallest of small-scale operations. Address and name management software is essential.

With many more addresses available, it is likely that the number of DNS entries will go up sharply as IPv6 is adopted. That, plus the physically larger addresses, means that more space will be needed for caches, for zone files, for support databases and so on. Nameservers will need more memory. If IPSEC is used, the computational cost of the DNS will rise, meaning that nameservers will need to be more powerful.

There is an important difference between **IPv6 aware** and **IPv6 capable**. A nameserver is **IPv6 aware** if it can answer queries about IPv6 addresses and resolve names to IPv6 addresses. A nameserver is **IPv6 capable** if it can be accessed via IPv6 itself.

Nameserver software is not really a problem. The reference DNS implementation and most widely-used nameserver, BIND, is already IPv6 aware and IPv6 capable. Other IPv6 aware and IPv6 capable nameserver software is also available, both open source and commercial. Most such software is also able to be run 'dual-stack', meaning it is able to answer queries that arrive via IPv4 or IPv6. Provided it is run on machines with sufficient resources, nameserver software should not be an obstacle to implementing IPv6.

The fact that IPv6 addresses are larger and more complex means that software to manage DNS information is almost essential. The days of hand-editing zone files are well over for IPv4, and will probably never be seen outside the testing lab for IPv6.

The extent to which existing in-house or third-party software can be adapted for IPv6 will be a major factor in the speed and effectiveness of IPv6 uptake. The larger the organisation, the more important this will be, and many smaller organisations that got by with IPv4 and partly manual systems will be driven to find suitable management software.

If IPSEC is used, or any of the IPv6 security features, then a Public Key Infrastructure is required (see below); an entirely new demand on management software and one that is very difficult to meet half-heartedly.

The benefit of the need for management software is that software-controlled DNS tends to be more reliable, less volatile and more stable than one that is updated by hand. Also, the need for such software creates a new market.

While much is said about the autoconfiguration facility in IPv6, it is likely that many sites will continue to need the control and flexibility offered by DHCP, in particular the facility offered by most DHCPv4 servers to perform proxy Dynamic DNS updates.

Unfortunately, there are no adequate DHCPv6 servers in existence, though there are several servers that are sufficiently advanced for testing or for smaller or less volatile networks. The lack of such software is a serious problem, and is currently delaying the uptake of IPv6 outside the lab.

For the DNS to work with IPv6, IPv6 network connectivity must be present. If a name resolves to an IPv6 address, and IPv6 connectivity is not present, then the remote host is not accessible. Since IPv6 queries resolve before IPv4 queries for the same name, the remote host is effectively inaccessible *even if* it has an IPv4 address as well. Luckily, simple measures can be taken to ensure a smooth transition (see **Part 3. Transition Support** below).

Because the DNS is a distributed database, being locally IPv6 aware is not enough. For IPv6 in the DNS to be fully operational, the root nameservers must also be IPv6 aware. That is, able to store IPv6 information and answer queries about IPv6 addresses. For a pure IPv6 system, the root nameservers must also be IPv6 capable, however this is not essential as long as clients remain IPv4 capable (or dual stack), which is likely to remain the case for some time to come.

The same is true of a local nameserver. If a nameserver is to answer queries from the Internet at large, it will need to be IPv4 capable, otherwise only IPv6 capable clients will be able to access it.

Because IPv6 is different in many ways to IPv4, even though the basic DNS operations are identical, a significant amount of training will be needed for system administrators, technical support personnel and others. Troubleshooting an encrypted, DNSSEC-secured IPv6 DNS transaction is significantly more complicated than troubleshooting an unsecured, open IPv4 DNS problem.

2.2 Summary

For hosts to be locatable via the DNS, the root nameservers need to be IPv6 aware. Many root nameservers and top-level domain nameservers are already IPv6 aware and can store IPv6 information. However, some IPv6 aware nameservers still do not yet return IPv6 answers [6]. Pure IPv6 hosts are not visible through these nameservers. Services to be accessed from the global Internet will need to be offered on dual-stack hosts at least until this changes.

Most of the problems with IPv6 that relate to the DNS can be solved with bigger, better hardware. Two that cannot be solved with hardware are the skills problem, and address and name management software. For organisations of all sizes, but particularly larger ones, the need to automate the flow of information from its source into the DNS will be a critical factor in how - or even whether - to move to IPv6.

2.3 Next Steps

An organisation wishing to implement an IPv6 aware DNS can do so now, with a few important caveats.

Any in-house network management or address management software needs to be inspected closely to see whether and to what extent it needs adapting to an IPv6 environment.

Important third-party network management or address management software in use in the organisation needs to be inspected equally closely, to ensure that it too will continue to operate in an IPv6 environment.

If DHCP and/or Dynamic DNS are in use in the organisation, thought will have to be given to how that functionality can be obtained for IPv6. If the systems managing DHCP/DDNS are manual or largely so, or if the network is highly homogeneous or very stable, then there are several suitable DHCPv6 servers that can be used as stopgaps. Otherwise implementation should concentrate on other aspects of implementation (testing, awareness building, education and planning for example), while suitable server software is developed.

Training of technical personnel is essential. Suitable programs should be developed early.

Part 3. Transition Support

3.1 Discussion

In this section, we discuss issues that will arise while transitioning to IPv6. Once an organisation has successfully moved to IPv6, these issues will lessen in importance or even disappear altogether. For a useful guide to deployment, see the Transition Checklist.

For reasons mentioned in **Part 2. DNS Operations** above, services available via IPv4 and IPv6 may need to be given two names during the transition phase - a name that maps to an IPv6 address and a name that maps to an IPv4 address.

Typically the existing name would be mapped to an IPv4 address and a new name mapped to the IPv6 address. For example, the server now called **service.ourdomain.com.au** might be given the new name **service6.ourdomain.com.au**, with the new name mapping to the IPv6 address of the service.

People with IPv6 connectivity can access the service via the new name, people without can access the service via the existing name. Once IPv6 is established, the names can be merged into one.

Connectivity during the transition phase will most likely be via tunnels of some sort, especially in cases where the provider does not offer IPv6 connectivity. Tunnels add some security issues, but do allow complete IPv6 connectivity. Larger organisations - those with high-end routing equipment and their own network management personnel - can craft their own tunnels, but for smaller organisations some kind of tunnel appliance will be needed, as tunnel-building requires significant network management skills. Some IPv6 tunnel appliances are already available [7], and others are in development.

Tunnelling (especially automatic tunnelling) poses security issues that native IPv6 doesn't have. Some good work has already been done on identifying transitional security issues [8].

Most high-end hardware is already IPv6 capable. Sadly most SOHO-level hardware (Small Office/Home Office) hardware is not, nor are most small CPE devices - the 'little black boxes' sold to small office and home users to connect them to ADSL and cable Internet providers.

The lack of suitable CPE devices is blocking widespread home deployment of IPv6, because ISPs cannot offer IPv6 to the home market, which is typically their bread and butter. Small office and home users wishing to move to IPv6 find themselves blocked as a result.

Cheap tunnel appliances get around this problem, but even very cheap devices cost money and the home market is already highly price competitive. The use of tunnel appliances also removes any pressure on ISPs to provide IPv6 connectivity.

On the other hand, as large markets like South-East Asia and China move towards IPv6 we should see commodity CPE devices being developed, and possibly even commodity tunnel appliances. Microsoft's newest operating system, Windows Vista, uses IPv6 by default. This will probably do much to spark home and SOHO interest in IPv6.

The biggest problems during transition will be skill deficiencies and software issues. Troubleshooting IPv6 at any level is significantly different to troubleshooting IPv4, and personnel at many levels will need some degree of re-training or further education, particularly system administrators and network support personnel. Incomplete or defective knowledge of IPv6 can cause damage directly (by doing something wrong) and indirectly (by not taking full advantage of the features offered by the new protocol).

Management will also need some level of education, as many aspects of transition require explicit support from the highest levels if they are to be successful. The business risks and the potential benefits need to be well understood if an organisation is to make the most of moving to IPv6.

Software developed in-house may need to be adapted to IPv6. This will of course depend on the particular software. Software involved in DNS and network management is at one end of the scale, and is likely to need substantial adapting. Well-written software, where a clear separation between application and network layers has been maintained, will be least likely to cause problems.

As noted in **Part 2. DNS Operations** above, managing the DNS with IPv6 is not really feasible without supporting software. Many larger organisations already have such software, and adapting it to IPv6 - especially if it was not designed with IPv6-readiness in mind - is likely to be a major task. Smaller organisations, beginning afresh, may find this aspect of transition easier.

Certainly *all* organisations, of any size, considering the development or purchase of software, should be considering IPv6 and ensuring that all specifications include IPv6.

Most third-party applications, and any in-house applications not directly concerned with naming and addressing, should be 'network agnostic' and are unlikely to pose significant problems. In general, any application that does not itself perform DNS lookups, but instead uses the facilities of the host operating system, will automatically support IPv6 as soon as the host operating system does. Operating system support for IPv6 is thus critical. Happily, the major operating systems are already largely IPv6 ready.

3.2 Summary

In-house network and DNS management software is likely to pose a major problem for larger organisations wanting to move to IPv6.

The lack of suitable CPE devices is blocking transition in the home and SOHO arenas, and thus the uptake of IPv6 by ISPs. Relief may come as large Asian markets start moving to IPv6. Tunnel appliances are a good stop-gap solution.

Third-party software is unlikely to be particularly problematical, as long as operating system support for IPv6 is present. In general, however, any software that stores, transmits, displays or uses IP addresses should be checked to make sure that the larger IPv6 addresses are properly handled.

3.3 Next Steps

Transition to IPv6 will look daunting because there is a lot of detail; however, it is possible now to implement IPv6 with minimal dificulty, provided good preparation is made for the transition. The Transition Checklist [9] is a good start.

Education, training and awareness programs are essential if transition is to be successful.

In-house and third-party software (including operating system software) will be the biggest item for most organisations. It should be checked carefully for IPv6 readiness, and plans made to upgrade, replace or adapt it as needed.

The remaining transition issues, particularly connectivity, are technical problems that have currently-available technical solutions.

Part 4. E-Commerce

4.1 Discussion

Organisations engaged in e-commerce and wishing to move to IPv6 will be dealing with all of the same issues that any other kind of organisation would have to deal with. This section deals with those additional issues that are more or less specific to e-commerce.

Some aspects of e-commerce are carried out on a very wide scale. Consider the number of deployed ATMs, EFTPOS and credit card readers - these devices make up a huge hardware base.

On one hand, IPv6 should be very interesting to those running very widely deployed networks and networks with numerous devices. On the other hand, getting all of those devices IPv6-ready - whether by upgrading or replacing them - is a major challenge. Getting it wrong will result in there being a lot of faulty devices to replace or upgrade again.

Such networks typically involve not only the software on a great many networked devices, but also some very large central programs. In a sense, such systems *consist* of software. Making sure this software is IPv6 ready is not only a big task, it is also a task where errors will be absolutely unacceptable. This factor will almost certainly cause reluctance to move to IPv6 unless there are very clear and compelling business reasons for doing so. Testing is likely to be extremely thorough, making the transition somewhat longer.

On the face of it there would seem to be excellent reasons to move from IPv4 to IPv6 for such widely deployed systems. IPv6 offers end-to-end security and payload encryption (where the content being carried is encrypted). IPv6's end-to-end transparency simplifies communications. And the huge address space not only simplifies addressing a large number of networked devices but also opens the way for entirely new classes of application based on massive deployment (RFID, sensor networks etc).

In the light of these technical advantages and the competitive advantages discussed below, IPv6 seems well worth adopting, even if making changes to a widely deployed platform does involve significant effort.

It seems likely that organisations moving to IPv6 will take full advantage of the security features of IPv6. This is especially likely in trading communities that can set up their own PKI. Organisations that do not adopt IPv6 risk being unable to participate in such markets.

IPv6, once implemented, will be an innovation-rich environment, just as IPv4 was and is. It is to be expected that new and effective ways will be found to use its features, especially the very large address space. Organisations not using IPv6 will be unable to participate.

E-commerce businesses using IPv6 will enjoy greater efficiencies and greater freedom to innovate. They will, in short, be more competitive. Organisations not moving to IPv6 will be at a competitive disadvantage.

In case the last few points seem overly hypothetical, it should be noted that our major trading partners, particularly Japan and South Korea, are already moving decisively towards IPv6.

4.2 Summary

The potential benefits to e-commerce from adopting IPv6 are great, but the transition is one that needs to be planned with the utmost care.

Software readiness is the biggest single issue for e-commerce. For e-commerce systems involving many deployed network devices such as credit-card readers, hardware readiness will also be a major issue.

The risk of lock-out is substantial if major trading partners move to IPv6, particularly if IPv6 security is then mandated within that trading community. The risk of becoming uncompetitive or of being left behind exists even if IPv6 security is not mandated.

4.3 Next Steps

Because of the high impact errors may have, the relatively long lead time for change, and the risk of lock-out, organisations engaged in e-commerce organisations should begin planning for IPv6 sooner rather than later.

Part 5. Security Keys

5.1 Discussion

IPv6 brings a new world of security features. The biggest are end-to-end trust and payload encryption. End-to-end trust means that it is possible for the recipient of a packet to know that the sender of that packet is the expected sender and that the packet contents received are the contents that were sent.

Payload encryption means that the contents of a packet cannot be eavesdropped in transit.

It is important to note here that the security features in IPv6 (IPSEC) secure the protocol, protecting the data in transit from alteration or theft. However, the protocol cannot identify, authenticate or authorise the participants that are sending and receiving that data. Higher level security, for example in the applications themselves, will still be necessary in many cases.

While the new features offer huge potention for more secure communications, they do require considerable new infrastructure developments, the greatest of which is a Public Key Infrastructure or PKI. A PKI is a support mechanism for creating, distributing, storing and revoking the keys used by the security features in IPv6.

Closed or restricted groups can set up their own PKI, but for Internet-wide security, an Internet-wide PKI is needed. It seems likely that IPSEC will initially be largely the province of closed user groups. This is no barrier to the uptake of IPv6 as such, but it is a serious obstacle to the widespread use of IPv6 security between arbitrary parties on the worldwide Internet.

There are a number of areas where IPv6 can be secured. For example, handshake communications between routers and hosts can be secured - i.e. router advertisments, neighbour discovery. Also, some existing infrastructure is used in new ways by IPv6 - the DNS is used by IPSEC, and reverse lookups take on greater importance.

Wherever security is implemented, cryptographic keys and hashes need to be generated and processed. Such things are computationally expensive. Secure IPv6 thus requires more processing power than IPv4. Current hardware is more than adequate for the task, but older, slower hardware, assuming its software can be upgraded, may suffer somewhat.

Security is a fairly complicated subject; implementing security badly can sometimes do more harm than implementing no security at all. For this reason, proper education, training and awareness programs are essential if IPv6 security features are to be used.

NAT (network address translation) was developed some years ago to deal with the fact that IPv4 addresses were running out, and IPv6 was (at the time) not ready for implementation. NAT is not needed with IPv6, and in fact actively interferes with IPv6 security, as well as preventing end-to-end transparency, one of IPv6's main advantages.

This concerns some network administrators, who see in NAT a form of primitive firewall. In fact, NAT is not, and has never been, an effective security tool. All the security functionality in NAT can be had far more cheaply and effectively in any simple packet filter (a packet filter is a subfunction of any firewall). A packet filter has the advantage that it is explicitly a security tool, rather than a non-configurable side-effect of a now-unneeded function.

5.2 Summary

The need for a PKI is probably the biggest obstacle to widespread uptake of IPv6 security. Closed or restricted communities (such as organisational intranets, trading communities and so on) can set up private PKIs, but until some kind of global PKI is present, IPv6 security will not be available to arbitrary correspondents on the global Internet. However, IPv6 can be implemented without the security features.

5.3 Next Steps

Understanding the benefits of IPv6 security is vital, as is setting those benefits against the costs and the requirements. It is important for the business case to be analysed, even if IPv6 security is then dismissed or postponed. The benefits are likely to be compelling, and should not be discarded lightly.

Conclusions

Organisations wishing to implement IPv6 can do so now. Depending on the nature of the enterprise, there are various caveats.

The major issue common to all implementations is the matter of in-house software, which may need to be upgraded, adapted or replaced.

A second major issue is third-party software; essential third-party software used in the enterprise (including operating system software) needs to be inspected for IPv6 readiness and if necessary replaced or upgraded. If IPv6 figures even to the smallest degree in future plans, new IT procurements should mandate IPv6 readiness.

The third common issue is education - education, training and awareness. IPv6 has enormous potential, but the enterprise may miss out on these, or even suffer damage, if those implementing IPv6 do so poorly or wrongly. This is especially true for IPv6 security.

For SOHO and home users, and for their symbiotic partners, Internet service providers, the single greatest barrier to IPv6 uptake is the lack of suitable CPE devices. This can be worked around in the short term using tunnel appliances, but these increase the cost of connectivity in a market which is already highly commoditised. Larger IPv6 SOHO and home markets in Asia may drive the development of more suitable low-cost CPE devices.

For large organisations or those with very volatile networks, lack of industrial strength DHCPv6 (and to a lesser extent DDNS) is a serious issue. Implementation can proceed using emerging DHCPv6 server software, but better products are needed.

References

[1] Within an organisation, IPv6 addresses can be used that are not routed to the global Internet. Such 'private' addresses need not be globally unique. Any host that is to be visible to the world must, however, have a globally unique address.

[2] See the IANA website - http://www.iana.org - and any RIR website for more on policy.

[3] NIRs and LIRs typically pay a membership fee to their parent body, which is proportional to the address space they are allocated. End-users of addresses, organisations using assigned addresses in operational networks, do not typically pay for their assigned address space.

[4] For example, APNIC - http://www.apnic.net/docs/policy/discussions/prop-035-v002.txt

[5] IP Version 6 Addressing Architecture - http://tools.ietf.org/html/rfc4291

[6] In technical terms, IPv6 'glue records' have not yet been configured on these servers. The .au nameservers are a case in point.

[7] Such as the 'Easy Access Device' being developed by the IPv6 for e-Business project. See http://www.ipv6.org.au/enable.html.

[8] For example, the draft IETF document on the subject. See draft-ietf-v6ops-security-overview-06.txt (a work in progress - '06' at time of writing).

[9] IPv for e-Business website - http://www.ipv6.org.au/transition.html

Glossary

BIND - Berkeley Internet Name Domain (previously Berkeley Internet Name Daemon) is the most commonly used name server on the Internet, especially on Unix-like systems, where it is a de facto standard.

CPE - customer-premises equipment, any terminal equipment and inside wiring located at a subscriber's premises and connected with a carrier's telecommunication channels, e.g. telephones, DSL modems or cable modems, or set-top boxes connecting to Internet Service Providers.

DDNS - Dynamic DNS, a system which allows the domain name data held in a name server to be updated in real time, e.g. to allow an Internet domain name to be assigned to a computer with a varying (dynamic) IP address.

DHCP - Dynamic Host Configuration Protocol, a set of rules used by a computer, router or network device to allow it to request and obtain an IP address from a server which has a list of addresses available for assignment.

DNS - Domain Name System, stores and associates many types of information with domain names, but most importantly, it translates domain names (computer hostnames) to IP addresses and IP addresses to domain names. In providing a worldwide keyword-based redirection service, DNS is a fundamental component of contemporary Internet use.

DNSSEC - Domain Name System Security Extensions, a suite of IETF specifications for securing DNS information, providing authentication and data integrity.

IPSEC - IP security - a suite of protocols for securing Internet Protocol (IP) communications by authenticating and/or encrypting each IP packet in a data stream. IPsec also includes protocols for cryptographic key establishment.

NAT - network address translation, involves re-writing the source and/or destination addresses of IP packets as they pass through a router or firewall. NAT exists solely as a workaround to allow one IPv4 address to be used by many hosts, at the cost of having those hosts "hidden" behind the NAT device, i.e. at the cost of losing end-to-end transparency, not to mention some efficiency, as every packet must be modified in transit.

PKI - Public Key Infrastructure, an arrangement that provides for trusted third party vetting of, and vouching for, user identities. This is usually carried out by software at a central location together with other coordinated software at distributed locations.

Zone files are nameserver configuration files. They contain information that defines mappings between domain names and IP addresses, and can also contain reverse mappings which resolve IP addresses into domain names.

Karl Auer
Internet Society of Australia and
Nullarbor Consulting Pty Ltd

The IPv6 for e-Business project is supported by the Australian Government through the Information Technology Online (ITOL) Program of the Department of Communications, Information Technology and the Arts.

[This document last modified Monday, 22-Jan-2007 11:30:22 EST]