Title:	**IPv6 Task Force**
	Mobile Wireless Working Group Report
	[Editor: B.Fernandes]
Author(s):	**IPv6 TF, Working Group 2**
Date:	**05.02.02**
Version:	**v1.3**
Distribution:	**Participants list, IPv6 TF lists**
Document number:	**74**

Table of Contents

1. The Business Model for IPv6

Given the remarkable growth of the Internet, and business opportunity represented by the Internet, IPv6 is of major interest to businesses, enterprise internetworks, and the global Internet. IPv6 presents all networking interests with an opportunity for global improvements, which is now receiving the collective action that is needed to realize the benefits. It is generally accepted that the first major commercial installation of IPv6 gear will be in upcoming 3G Wireless Networks. As wireless operators plan for the deployment of millions of IP devices (phones, cars, PDAs, home appliances...) in the 3rd generation networks, the shortage of IPv4 address space is a major impediment in Europe, Asia and South America. By providing globally unique, stable addresses for all Internet devices, IPv6 is the solution to accommodate the estimated 1 billion subscribers in 2002, and the 24 trillion bits of data that will be exchanged through 3G networks in 2004.

IPv4 has proven to be robust, easily implemented and interoperable, and has stood the test of scaling an internetwork (a network of networks) to a global utility the size of today's Internet. While this is a tribute to its initial design, moving forward to an even grander scale requires laying a new foundation. IPv6 is not a simple derivative of IPv4, but a definitive improvement. Auto-configuration, security, and support of real-time communications could only be partially integrated in IPv4. They are mandatory in all implementations of IPv6.

While the Fixed wired line networks are already using IP and Session Initiation (SIP) to offer Multimedia Services, 3GPP the Cellular standards bodies has chosen SIP and IPv6 to be mandatory protocols for its IP Multi-media Subsystem (IMS).

IP based Multi-media services will be launched on 3G networks in order to enable operators to move smoothly towards the full-IP multimedia target. These will be supported using functionality known as the IP Multi-media Sub-system (IMS) which is being specified as part of 3GPP Release 5, currently planned for completion in December 2001 and released in ~ March of 2002 with possible subsequent functionality releases in October 2002. The IMS will make exclusive use of IPv6 and is being designed (initially) to operate over GPRS packet transport capabilities. The reasons for adopting IPv6 exclusively within the IMS include the “always on” connectivity, ready to transmit and process information at any time without delays. Besides combined with the rapid growth in cellular devices and the potential lack of public IPv4 addresses makes it inevitable for this paradigm to take place without IPv6. Additionally, adopting IPv6 support in the future devices to be in a state of permanent standby, IMS will open a significant number of new business opportunities and is expected to be an important part of future mobile networks. It represents one of the key UMTS features by which operators will be able to differentiate themselves from competitors. In addition, for UMTS Operators it represents a key differentiator from present GSM 2G and 2.5G networks. Consequently, it should be introduced as early as possible in order to maximise service opportunities and customer benefits.

Mobile operators are in the process of introducing packet services using GPRS – an IP based solution for GSM and 3G UMTS networks which is being initially implemented using IPv4. It should be noted, however, that the supporting GPRS network could be either IPv4 or IPv6 based.

2. The Need for further Development

Adding new elements of Internet technology into the wireless domain, in addition to the growth potential of wireless always on/always with you clients will gradually require a smooth migration towards IPv6. Combination of VOIP, wireless LANs, SIP and ENUM, all of which call for further developments in a future communications environment, will require IPv6 in an early deployment

phase giving Operators / Carriers a leading edge of further developing their potential business case.

- With IPv6 everywhere, mobile users can get a seamless Internet experience. 3G providers will be able to offer the full range of Internet capabilities, in addition to their existing voice/telephony services. Users can connect to whatever web sites they choose, log in to their corporate Intranet (and be reached from that network), do VOIP, get streaming audio/video, use whatever network applications they need, etc. They are not constrained to the (limited set of?) value added network services the 3G operators offer: e.g. second generation WAP gateways.
- Addressing and naming for the operator becomes easier. There will be no address space exhaustion and no need to resort to messy and non-scalable workarounds like schemes based on NAT and on-the-fly IPv4/IPv6 mapping.
- Mergers and acquisitions by 3G operators are simplified to the point of becoming trivial, at least in networking terms. Each 3G operator would presumably have its own set of globally unique IPv6 address space (i.e. not using 10/8 in the IPv4 world like GSM operators today offering GPRS). When one 3G company merges with another, their networks can just be bolted together and everything will just work automatically. There's no address migration to do or renumbering or NAT deployment.
- Having a homogeneous network will reduce network management costs. E.g., Routers and network management stations/systems don't have to deal with IPv4 and IPv6 packets. Provisioning the network's naming and address databases is simpler: there's only one addressing scheme to maintain. Likewise, packet routing is simplified. The same goes for the DNS: an all IPv6 approach is much cleaner and cheaper to set up and manage than a combined IP v4/IPv6 one.
- In the initial phase of GPRS/UMTS, IPv4 is a perfectly reasonable solution. Using NAT‘s, IPv4 and private address space will of course offer a market service for the early adopters. However to reach full interoperable services and an unlimited scalable solution, IPv6 is a key facilitator. Therefore a transition to IPv6 is necessary. From a smooth transition and investment perspective, the sooner the better, however there are no definite limits, but a gradual increase without a unique breakpoint.
- There are severe problems and limitations with band-aids such as NAT, and although these band-aids and extensions may prove valuable in the near term, they ultimately will limit connectivity, interoperability, and performance in enterprises that are increasingly network-dependent.
- While end-user and business requirements for advanced network services expand exponentially, a protocol such as IPv4 will not be able to cope.
- There will not be a magic date imposed on anyone to move to IPv6, but rather a simply incentive to move before it becomes too late and too expensive.
- The coexistence of IPv4 and IPv6 will last many-many years and that the phasing out of IPv4 will be soft and gradual, (there are still a few thousand users of analogue phones in the Nordic countries).
- Ultimately, there is no substitute for IPv6 whenever multimedia, interactive, and transaction-oriented network applications require high levels of connectivity. IPv6 will become critical to the operations and continued efficiency of day-to-day business activities.
- IPv6 will create new service opportunities and customer benefits for Cellular Network. Besides IPv6 will bring:

  - Operational benefits, including network efficiency.

  - Cost efficiency (i.e. potential reduction of future operating costs).

- Minimising industry-wide disruption.

In essence, the Mobile community has a unique chance to investigate and pioneer the future, together with all other Internet related players, be they fixed, cable, ADSL, ISPs etc. In so doing they will acquire a competitive edge which can be exported.

3. Choice of IP version

The GPRS and UMTS standards allow a choice of IP version - IPv4 or IPv6.

All GPRS networks today use IPv4. There are many reasons for this with the most significant since all available infrastructure and terminals are based on IPv4. Although IPv6 appears on vendor's product roadmaps and the first IPv6 –capable solution was announced by in November 2000 it is unlikely that all infrastructure and terminal suppliers will have even experimental products available before 2002.

Given the amount of emphasis by networks operators and infrastructure vendors on UMTS implementation – particularly with the very high license fees being paid by European operators – what was questionable whether the implementation of IPv6 on GPRS will ever be justifiable is no more valid and some vendors will have implemented IPv6 in GPRS.

There are three potential drivers for moving to IPv6:

- Meeting the requirements in the standards, for example IPv6 is required for IP multimedia (although only for the IMS domain)
- Avoiding problems when IPv4 reaches exhaustion
- Obtaining benefits from features that IPv6 offers that are not available in IPv4.

A significant uncertainty is the speed with which IPv6 will be introduced generally in the Internet world. Here there are two extremes and a continuum of possibilities between them.

- The first extreme is that ISPs will perceive some real operational advantage in using IPv6 and will introduce it as soon as possible in order to capitalise on these advantages.
- The other extreme is that Network Operators / ISPs will regard the introduction of IPv6 as an avoidable expense and will delay its introduction as long as possible, i.e. until the shortage in IPv4 addresses begins to be felt.

UMTS Release 3 offers options to use either IPv4 or IPv6 (see 3G TS 23.003 V3.5.0 June 2000 section 3.7 and 3.8) for mobile terminals and specifies that either permanent or temporary allocation may be used. In practice, they will need to support IPv4 for general compatibility and may choose to support IPv6 as well with a dual stack.

UMTS Release 3 TS 23.003 specifies that the GSN address may use either IPv4 or IPv6, but the GTP specifies that IPv4 is mandatory and that IPv6 is an optional addition.

UMTS Release 5 specifies in 3G TR 23.821 V1.0.1 (2000-07)[1] section 11.1 that:

- Network elements of the IP Connectivity services (between RNC, SGSN and GGSN) and IP transport for the CS Domain may continue to use either IPv4 or IPv6.
- Terminals shall be able to access data services based on IPv4 and IPv6

[1] The latest version of the 3GPP specifications can always be found on the 3GPP Servers and valide.

- Network elements in the IP multimedia domain shall be based exclusively on IPv6

Phase	SGSNs GGSNs etc.	Mobile terminals and external services	Network elements for IP multimedia
GPRS Release 98	IPv4	IPv4 or IPv6 In practice it will be IPv4, IPv6 optional	
UMTS Release 3,4	IPv4, IPv6 optional	IPv4 or IPv6 In practice it will be IPv4, IPv6 optional	
UMTS Release 5	Not yet decided, Possibly the same as R99	IPv4 and IPv6 (IPv6 exclusively for multimedia)	IPv6 exclusively

For the core network prior to UMTS Release 5, the main issue is the support of the GPRS tunnelling protocol. The specifications concerned are:

- GPRS Release 98: EN 301 347 (GSM 09.60)
- UMTS Release 3; 3G TS 29.060
- The equivalent specification for Release 5 can be found on the 3GPP Server.

The specifications for the GTP from release 98 onwards accommodate both IPv4 and IPv6 as addresses for terminals. These addresses pas transparently through the tunnel and therefore the version used by the terminal and service is independent of the version used by the core network in support of the tunnel.

Operators are therefore faced with decisions about when to introduce IPv6 into the core network and start dual stack working and whether to do so in advance of needing IPv6 for IP Multimedia. This decision will be influenced by their view on IPv4 exhaustion and feature advantages of IPv6.

4. Feature advantages of IPv6

The following are the main advantages of IPv6 in terms of features: Conservation: IPv6 will restore the paradigm of end-to-end functionality. This was disrupted by NAT (Network Address Translation), where the address of a packet from the internal network (mostly using private addresses, not routable in the Internet) has to be exchanged for an official IP address to be routed on Internet. This is decreasing performance as every packet has to be analysed and its header changed. This breaks checksums, end-to-end Security and applications, that needs a fixed IP address such as some forms of IP Telephony.

Routing: IPv6 uses a routing hierarchy with aggregation. In the old days of the Internet the address spaces have been distributed all over the world without any really idea how the routing ca be constructed based on this distribution. This changed with the introduction of CIDR (Classless InterDomain Routing), and the registries created allocation policies supporting route aggregation and appropriately sized address ranges.

This has slowed down the dynamic growth of the routing tables in the whole Internet but it cannot reverse the mistakes of the first the years and the large and unstable routing tables that remain.

Renumbering is difficult and changing the provider either involves renumbering or creates holes in the aggregation block of the old provider, increasing the routing tables. IPv6 will provide hierarchical address allocation limiting the number of entries in the routing tables to about 8000 compared to some 90,000 with IPv4 at present.

Plug & Play: IPv6 reduces the administration and management overhead by making plug & play really work. Auto configuration works together with the Dynamical Host Configuration Protocol and the Domain Name System, so the system administrator is not forced to configure every workstation and PC manually. The address is a combination of a routing part (prefix, 64 bits) and a host ID (EUI-64, 64 bits). The auto configuration mechanism reads the MAC address and composes a network wide valid ID. The prefix is provided by a local facility and can be changed if necessary (change of provider) without difficulty and without manual configuration of the host. The cost savings in administration and management can be quite substantial.

Mobility: IPv6 supports mobility much better than IPv4. All IPv6 networks and nodes are ready for mobile IPv6. IPv6 Neighbour Discovery and Address Auto configuration allow hosts to operate in any location without a special support. The performance is improved because of traffic optimisation. The flexible address structure is well suited for roaming. Extended security concept might be adopted to meet the higher requirements from the mobile world.

Header structure: IPv6 has an optimised header structure. Unlike IPv4, the header of IPv6 has fixed sizes and fewer fields. This will make faster processing possible, and implementations in hardware will give the needed performance for fast networks. The option fields will be processed only if the option is present. However, first implementations will not operate IPv6 in hardware.

Security: IPv6 will provide means for privacy and security as an integral part of the standard rather than as a separate protocol. With IPv4, the IPsec protocol is used, which is not different in principle to IPv6, but is very complex and difficult to use. Before IPsec can be used in communication, it requires a check to see if the peer is supporting IPsec at all and what are the implemented features.

4.1 Issues to consider when introducing IPv6

"The transition to IPv6 is a non-trivial exercise… Making IPv4 and IPv6 systems interwork is equally challenging… The fastest path along that road involves serious hands-on experience with the protocols and applications that use it." Vinton Cerf, honorary chairman of the IPv6 Forum.

IPv6 should not be thought of as just another IP protocol, IPv6 is an entirely different protocol from IPv4. There is compatibility between IPv4 and IPv6, and inter working is only achieved by protocol translation devices or by enabling all devices to use both protocols. IPv6 is an immature technology that is still developing. There is very little experience of operating IPv6 in production networks, Ipv6 has not been extensively tested in a production environment and until this testing has been completed, we cannot be sure we have captured all the issues associated with using IPv6. There are many issues with IPv6, most of which are technically very detailed and there is no simple overarching architectural answer to them. A summary of just the performance, costs and development issues wit IPv6 follows:

Performance:

- Currently implementations of IPv6 on routers are slower than IPv4.
- Filtering packets based on IPv6 addresses and IP port numbers may not be possible because of the extra width of address fields that needs to be compared may exceed CAMs (Contents Addressable Memory) in the router hardware.
- The extra 24 bytes overhead IPv6 has above IPv4 per VoIP packet causes significant inefficiencies.

Costs:

- IPsec support is mandatory in IPv6 devices even in devices like routers, which do not strictly need the IPsec mechanism. IPsec requires PKI (Public Key Infrastructure) and CA (Certificate Authorities) to be operated in a scalable manner, these additional components incur extra capital, maintenance and management costs.
- IPv6 claims to do away with NAT but since IPv4 and IPv6 will co-exist for a long time then transition mechanisms, such as NAT-PT, will have to be maintained for a long time. Many of the transition mechanism, which translates between IPv6 and IPv4, are more complex than existing IPv4 NAT mechanisms and will therefore be more expensive to build and maintain.
- Since IPv4 and IPv6 will coexist for a long time many IP networks will have to run dual protocol, that is both IPv4 and IPv6. The addition of IPv6 will require additional configuration of the network. For large ISPs this amounts to hundreds of thousands of lines of configuration code.
- In Network Operation Centres the IP address of the customer is a basic piece of information that is used I the majority of support requests. Has their been any practical considerations of the increased time take to read out an IPv6 address accurately compared to the simple 4 decimal numbers of IPv4?
- IPv6 imposes an additional layer in the IP address allocation hierarchy. Currently for IPv4 RIRs (Regional Internet Registries) allocate IP addresses to LIRs (Local Internet Registries) which assign IP addresses to end-users. In IPv6 RIRs, allocate addresses to TLAs who allocate addresses to NLAs, which assign IP addresses to end-users. This extra depth of hierarchy imposes additional bureaucracy and requires additional IP address management components.
- IPv6 Auto-configuration only applies to end-stations and not routers. IPv6 actually imposes an additional configuration overhead on the router configurations.
- Fixing unknown bugs. Early users of IPv6 will experience many bugs; the fixing of these bugs will cost IPv6 customers and network providers.

  OSS & NMS development. Current vendors of IP network OSS and NMS although aware of IPv6 have not yet committed to redeveloping their applications to allow management of IPv6 networks. Current indications from OSS and NMS vendors are that the cost of these developments will be significant.

  IPv6 should be considered as a disruptive technology that today and for the next, few years will cost more, operate slower and have less functionality than IPv4. Eventually IPv6 will become the dominant IP technology when it has been through a development and customer acceptance cycle. The issue is when should UMTS use IPv6 and is it willing to use its green field networks as an IPv6 test bed?

4.2 IPv6: Standardization and Production Status

IPv6, the Next-Generation Internet Protocol, has been approved as a Draft Standard, so that it is known to be highly stable and appropriate for production. A large number of end-user organizations, standards groups, and network vendors have been working together on the specification and testing of early IPv6 implementations. A number of IETF working groups have produced IPv6 specifications that are finished or well underway. Standards work on IPv6 and related components is far enough along that vendors have already committed to a considerable

number of development and testing projects. All of the major router vendors have made plans to support IPv6 in their products. Cellular Standards bodies are offering the industry interoperability test facilities, training programs and the IPv6TF offers the possibility of participating in Trails.

Furthermore, Network software vendors have announced a wide range of support for IPv6 in network applications and communication software products. Software is available from Microsoft for Windows-based clients. Microsoft is actively working on support for the emerging update to the Internet Protocol, IPv6 because IPv6 brings back the capability of 'end-to-end control of communications'; making networking applications simpler as the network again becomes transparent.

Microsoft .NET will allow the creation of truly distributed services that will integrate and collaborate with a range of complementary services to serve customers in ways that can only be dreamt of today. Achieving this level of service distribution will require the interconnecting Internet infrastructure to provide transparent and consistent views of the environment to the participating systems, applications, and services.

For example, as the popularity of multi-player games conducted over the Internet increases, the need for consistent network layer addressing between the participating peers does as well. IPv6 enables consistent views because there is no need for the address sharing which results in peers having differing views of the address for a participating system. In addition to a consistent view of addresses, IPv6 allows multiple machines within a home to use the same port number (many services use a registered port number for inbound connections), where the currently prevalent IPv4 NATs restrict port number use to one machine at a time. Removing the possibly of interfering technologies from the environment frees the game developer and player to focus on the game play.

Additional peer-to-peer applications that are made easier using IPv6 include IP telephony, and video tele-conferencing. These, and applications like them, are likely to take advantage of the Quality-of-Service (QoS) features defined for IPv6. While many of the QoS features have also been defined as add-ons for IPv4, the mechanism selected was to redefine the meaning for one field of the IP header, causing collisions with historical implementations. The effort to provide QoS for IPv4 has struggled due to differing models of deployment. This effort is not wasted though, because it is forcing many details to be worked through from hardware capabilities to business practices. IPv6-enabled systems will be able to leverage this effort to provide an array of service levels that are consistent from end-to-end.

Wireless technologies are emerging which present the prospect of ad-hoc networks between personal devices. Setting up systems to work in an ad-hoc mode is challenging enough, but it is expected that many of these personal devices will also need to work in the managed environment of the workplace. Switching between these modes is frequently frustrating and is significantly more involved than either alone. To avoid the complexity, IPv6 has defined an architectural principle that systems are required to simultaneously support multiple addresses. Coupling this capability with scoped addresses results in the ability to move easily and automatically between ad-hoc and managed environments.

Another capability IPv6 brings to the wireless realm is efficient mobility. Many applications today expect the IP addresses to remain constant throughout the lifetime of their connection to a remote peer or server. While this is possible today using IPv4, the mechanisms are complex, and operationally very fragile. IPv6 removes much of the complexity and allows the end systems to efficiently redirect packets to the new address of the mobile node via a binding update. By maintaining the awareness of mobility in the endpoints, the tremendous power of a flexible transparent network is preserved.

To make information available any time, any place and on any device, requires the constellations of devices to have a consistent view of each other, which will be possible using the globally unique addresses available with IPv6. With this capability, every application on any device can be

exposed as a service on the Internet. The programming model gives developers the opportunity to focus fewer resources on where or how an application runs over the potential network impediments like NAT, and more on what it does. Like the rest of the Microsoft .NET tools, the IPv6 implementation will automatically adapt itself to the current needs, be that ad-hoc, home, or business connections.

To address concerns about security and privacy, some vendors have in their implementation IP layer security (IPSec) included. IPSec is an industry standard security technology that provides for data authenticity and integrity as well as privacy of communications across the array of protocols used by the various applications. Providing the capability at the network layer frees the developer from having to add specific security capabilities to every application.

To make stateless auto configuration work well, the standards community chose the underlying hardware address for use as part of an IPv6 address to ensure global uniqueness. This has the side effect that all communications are traceable to the specific hardware device. While it is technically necessary to have a published (over some scope), globally unique address to receive incoming connections, the address of an originator only requires current global uniqueness (not publication). The ability of IPv6 systems to simultaneously support multiple addresses allows each application to use an independent address, and/or an application to use a different address for each service to which it connects.

4.3 Getting There

It will affect nearly all networked applications, end-systems, infrastructure systems, and network architectures. It is critical that this change be approached with responsibility to prevent costly unproductive missteps that result from broad premature availability of technologies.

It is anticipated that Internet service providers (ISPs) will react to customer demand as the deciding factor for when to deploy native IPv6 routing, but as it takes several years to replace the network equipment this may be a slow process. To avoid a "chicken-and-egg problem", encapsulating IPv6 packets within IPv4 will allow incremental deployments of end systems that will in turn demonstrate the demand to the ISPs. To stay on the high performance path of the existing routers, IPv6-enabled Windows systems will default to tunnelling over IPv4 unless the ISP provides specific indication to do otherwise and a native IPv6 path exists end-to-end. The only requirement is that the Windows systems directly connected to an ISP receive at least one public IPv4 address (the address ranges specified in RFC 1918 are not public). Subsequent systems in a home or business will receive 6to4 prefix router advertisements from the directly connected system.

The next step will be to enable key components of the system for IPv6 so developers can begin the task of IPv6 enabling the applications. It is also expected that early adoption customers will start using IPv6 under non-production, controlled conditions. This will allow those customers to have better visibility into managing their eventual rollout, and it will help identify any issues that need to be addressed in networked products.

4.4 IPv6 Market Impact

Despite the approach, the transition will not be easy. It is expected that for the foreseeable future, most manufacturers will produce systems supporting both IPv4 and IPv6, so that if connections are not possible using IPv6 they can fall back and succeed using IPv4 (providing IPv4 connectivity existed prior to the introduction of IPv6). The overall goal is to ensure a smooth transition and deployments where updated applications can take advantage of the new protocol, without breaking existing applications. To this end there have been new APIs defined to specifically isolate the legacy applications from unintentional exposure to protocol differences, including the larger IPv6 addresses.

4.5 Why IPv6?

IPv6 is the next generation Internet Protocol designed as a successor to the current version IPv4. IPv6 enables a high performance, scalable Internet. The specification includes a number of enhancements over IPv4, including:

The Internet originated as a research network connecting stationary computer equipment, used by a few for transporting primarily ASCII files. It has grown into an essential commercial and business tool, sometimes-running mission critical applications. It is evolving into a transport vehicle able to connect a multitude of handheld and wireless devices, supporting applications such as bank transactions and Internet Telephony. This next generation Internet will connect billions of devices, some with multiple IP addresses. Internet users will demand the global connectivity, transaction security, uncompromising performance, expandability and affordability. IPv6, the next generation Internet Protocol helps to meet this challenge. The core IPv6 protocol and many of the related RFCs are IETF draft standards.

IPv6 offers a near infinite number of globally unique IP addresses to enable easier implementation of peer-to-peer computing and "Push" applications. IPv6 mandates native authentication and security. It will also offer true mobility. The end user application will use flow-label management to provide Differentiated Services and Quality of Service features to real time applications like IP Telephony. Furthermore, IPv6 will reduce the administrative workload by enabling "plug-and-play" (automatic) address configurations. It will ease network addresses re-numbering requirements with built-in auto-configuration and hierarchical addressing schemes. It will eliminate the need for the Network Address Translation (NAT) function and associated IPv4 performance bottlenecks and application limitations.

The IPv6TF advocates a smooth migration, supporting coexistence of both IPv4 and IPv6 during the transition. This will enable customers to leverage their existing investment of today's IPv4 services, while preparing for a seamless migration to IPv6 as additional IPv6 devices come online. The Industry is encouraged to continue to aggressively bring the cost and performance benefits of emerging technologies, such as IPv6, online as standards based solutions.

4.6 IPv6 Design Goals

IPv6 has been designed to enable high performance, scalable internetworks that should operate as needed for decades. Part of the design process involved correcting the inadequacies of IPv4. IPv6 offers a number of enhanced features, such as a larger address space and improved packet formats. Scalable networking requires careful

Utilization of human resources as well as network resources; so, a great deal of attention has been given to creating auto configuration protocols for IPv6, minimizing the need for human intervention.

Other benefits relate to the fresh start that IPv6 gives to those who build and administer networks. For instance, a well-structured, efficient and adaptable routing hierarchy will be possible

4.7 Addressing and Routing

IPv6 helps to solve a number of problems that currently exist within and between enterprises. On the global scale, IPv6 will allow Internet backbone designers to create a flexible and expandable global routing hierarchy. The Internet backbone, where major enterprises and Internet Service Provider (ISP) networks come together, depends upon the maintenance of a hierarchical address system, similar to that of the national and international telephone systems. Large central-office phone switches, for instance, only need a three-digit national area code prefix to route a long-

distance telephone call toward the correct local exchange. The current IPv4 system also uses an address hierarchy to sort traffic towards networks attached to the Internet backbone.

Without an address hierarchy, backbone routers would be forced to store route table information on the reach ability of every network in the world. Given the current number of IP subnets in the world and the growth of the Internet, it is not feasible to manage route tables and updates for so many routers. With a hierarchy, backbone routers can use IP address prefixes to determine how traffic should be routed through the backbone.

Legacy IPv4 address assignments that originated before CIDR and the current access provider hierarchy often do not facilitate summarization. The lack of uniformity of the current hierarchical system, coupled with the rationing of IPv4 addresses, makes Internet addressing and routing quite complicated. These issues affect high-level service providers and consequently individual end users in all types of businesses. Furthermore, renumbering IPv4 sites when changing from one ISP to another, to maintain and improve address/route aggregation is unnecessarily complicated (and thus more expensive) compared to IPv6's ease of site renumbering.

Users in private address spaces with non-unique addresses typically require gateways, and possibly Network Address Translators (NATs), to manage their connectivity to the outside world. In such situations, some services are simply not available. A NAT is meant to allow an enterprise to have whatever internal address structure it desires, without concern for integrating internal addresses with the global Internet. This is seen as particularly convenient in the existing IPv4 world, with its more cumbersome address space management. The NAT device sits on the border between the enterprise and the Internet, converting private internal addresses to a smaller pool of globally unique addresses that are passed to the backbone and vice versa.

NAT Network Address Translator (NAT) may be appropriate in some organizations, particularly if full connectivity with the outside world is not desired. But for enterprises that require robust interaction with the Internet, NAT devices often get in the way. The NAT technique of substituting address fields in each and every packet that leaves and enters the enterprise is very demanding, and presents a bottleneck between the enterprise and the Internet. A NAT may keep up with address conversion in a small network, but as the enterprise's Internet access increases, the NAT's performance must increase in parallel. The bottleneck effect is exacerbated by the difficulty of integrating and synchronizing multiple NAT devices within a single enterprise. Enterprises with NAT are less likely to achieve the reliable high performance Internet connectivity that is common today with multiple routers attached to an ISP backbone in an arbitrary mesh fashion. Furthermore, use of NAT devices takes away the additional element of reliability afforded by the possibility for asymmetric routing, since NAT devices require control of traffic directions both to and from internally addressed network nodes.

NAT translators also run into trouble when applications embed IP addresses in the packet payload, above the network layer. This is the case for a number of applications, including certain File Transfer Protocol (FTP) programs, Mobile IP, and the Windows Internet Name Service (WINS) registration process of Windows 95 and Windows NT. Unless a NAT parses every packet all the way to the application level, it is likely to fail to translate some embedded addresses, which will lead to application failures. NAT can also break Domain Name Servers, because they work above the network layer. NATs prevent the use of IP-level security between the endpoints of a transaction. Today, NAT devices are helpful in certain limited scenarios for smaller enterprises, but are considered by many to be generally disadvantageous for the long-term health of the Internet.

IPv4's limitations also aggravate the occasional need in many organizations to renumber network devices -- i.e., assign new IP addresses to them. When an enterprise changes ISPs, it may have to either renumber all addresses to match the new ISP-assigned prefix, or implement Network Address Translation devices (NATs). Renumbering may be indicated when a corporation

undergoes a merger or an acquisition with consequent network consolidation. Since routing prefixes are assigned to reflect the routing topology of the enterprise networks and the number of nodes attached to the particular network links, there are two ways that the choice of routing prefixes can become inconvenient or incorrect:

1. The routing prefix can become too long for the administration to be able to increase the number of nodes that can be attached to the particular link, and

2. The ways that the network links are connected together, or are connected to the outside world can change.

Either of these occurrences would indicate the need to renumber one or more enterprise networks. It would be quite profitable to be able to renumber enterprise networks without requiring expensive downtime for the networks and or the nodes on the network.

Address shortages and routing hierarchy problems threaten the network operations of larger enterprises, but they also affect small sites -- even the home worker who dials in to the office via the Internet. Smaller networks can be completely dropped from Internet backbone route tables if they do not adapt to the address hierarchy, while larger networks may refuse to renumber and cause a larger routing problem for the backbone providers of the Internet. With today's IPv4, address registries, ISPs with individual dial-in clients cannot allocate IP numbers as freely as they wish. Consequently, many dial-in users must use an address allocated from a pool on a temporary basis. In other cases, small dial-in sites are forced to share a single IP address among multiple end systems.

A unique IP address sets the stage for users to gain direct connectivity to other users on the Internet, as determined by local policy. It also simplifies a wide range of productive interactive applications, of which telecommuting and remote diagnostics are only two examples. Today's hierarchy of limited and poorly allocated IPv4 addresses has already caused problems, and will continue to do so as more and more devices of varying capabilities are added to the Internet.

4.8 Security

Encryption, authentication, and data integrity safeguards are needed for enterprise internetworking and virtual private networks (VPNs). For these purposes, IPv6 offers security header extensions.

The IPv6 authentication extension header allows a receiver to determine with a high degree of certainty whether or not a packet originated from the host indicated in its source address. This prevents malicious users from configuring an IP host to impersonate another, to gain access to secure resources. Such source-address masquerading (spoofing) is among the techniques that could be used to obtain valuable financial and corporate data, or could give adversaries of the enterprise control of servers for malicious purposes. Spoofing might fool a server into granting access to valuable data, passwords, or network control utilities. IP spoofing is known to be one of the most common forms of denial-of-service attack; with IPv4 it is typically impossible for a server to determine whether packets are being received from the legitimate end node. Some enterprises have responded by installing firewalls, but these devices introduce a number of new problems, including performance bottlenecks, restrictive network policies, and limited connectivity to the Internet or even between divisions of the same company.

4.9 Mobility

IPv4 has difficulties managing mobile computers, for several reasons:

- A mobile computer needs to make use of a forwarding address at each new point of attachment to the Internet, and it's not always so easy to get such an address with IPv4

- Informing any agent in the routing infrastructure about the mobile node's new location requires good authentication facilities, which are not commonly deployed in IPv4 nodes.

  In IPv4, it may be difficult for mobile nodes to determine whether or not they are attached to the same network.

- It is unlikely in IPv4 that mobile nodes would be able to inform their communication partners about any change in location.

Each of these problems is solved in a natural way by using features in IPv6. The benefits for mobile computing are apparent in quite a number of aspects of the IPv6 protocol design, and go beyond merely providing dial-up support for road warriors. The improvements in option processing for destination options, autoconfiguration, routing headers, encapsulation, security, and anycast addresses all contribute to the natural design of mobility for IPv6 [19]. In fact, some satellite work in Europe is already starting to become IPv6 based. Combining flow label management to provide better Quality of Service to mobile nodes may further emphasize the IPv6 mobility advantage.

4.10 The IPv6 solution

IPv6, with its immensely larger address space, defines a multi-level hierarchical global routing architecture. Using CIDR-style prefixes, the IPv6 address space can be allocated in a way that facilitates route summarization, and controls expansion of route tables in backbone routers. The vastly greater availability of IPv6 addresses eliminates the need for private address spaces. ISPs will have enough addresses to allocate to smaller businesses and dial-in users that need globally unique addresses to fully exploit the Internet. Using an example from crowded telephone networks, one might say that IPv6 eliminates the need for "extensions", so that al offices have direct communication lines and do not need operators (automatic or otherwise) to redirect calls.

4.11 Address Auto configuration

Each IPv6 node initially creates a local IPv6 address for itself using "stateless" address auto configuration, not requiring a manually configured server. Stateless auto configuration further makes it possible for nodes to configure their own globally routable addresses in cooperation with a local IPv6 router. Typically, the node combines its 48 or 64 bit MAC (i.e., layer-2) address, assigned by the equipment manufacturer, with a network prefix it learns from a neighbouring router. This keeps end user costs down by not requiring

Knowledgeable staff to properly configure each workstation before it can be deployed. These costs are currently part of the initial equipment expense for almost all IPv4 computing platforms. With the possibility of low or zero administrative costs, and the possibility of extremely low cost network interfaces, new market possibilities can be created for control of embedded computer systems. This feature will also help when residential networks emerge as an important market segment.

IPv4 networks often employ the Dynamic Host Configuration Protocol (DHCP) to reduce the effort associated with manually assigning addresses to end nodes. DHCP is termed a "stateful" address configuration tool because it maintains static tables that determine which addresses are assigned to newly connected network nodes. A new version of DHCP has been developed for IPv6 to provide similar stateful address assignment as may be desired by many network administrators. DHCPv6 also assists with efficient reconfiguration in addition to initial address configuration, by using multicast from the DHCP server to any desired population of clients.

The robust auto configuration capabilities of IPv6 will benefit internetwork users at many levels. When an enterprise is forced to renumber because of an ISP change, IPv6 auto configuration will

allow hosts to be given new prefixes, without even requiring manual reconfiguration of workstations or DHCP clients. This function also assists enterprises in keeping up with dynamic end-user populations. Auto configuration allows mobile computers to receive valid forwarding addresses automatically, no matter where they connect to the network.

4.12 Multicast

Modern internetworks need to transmit streams of video, audio, animated graphics, news, financial, or other timely data to groups of functionally related but dispersed end stations. Network layer multicast best achieves this. Typically, a server sends out a single stream of multimedia or time-sensitive data to be received by subscribers. A multicast-capable network routes the server's packets to each subscriber in the multicast group using an efficient path replicating only as needed.

Multicast applications have been developed for IPv4, but IPv6 extends IP multicasting capabilities by defining a much larger multicast address space. All IPv6 hosts and routers are required to support multicast. In fact, IPv6 has no broadcast address as such; it has various multicast addresses of various scopes. The improvements offered in IPv6 promises to simplify the use and administration of multicast in many applications.

4.13 Anycast

Anycast services, supported in the IPv6 specification, are not defined architecturally in IPv4. Conceptually, anycast is a cross between unicast and multicast: an arbitrary collection of nodes may be designated as an anycast group. A packet addressed to the group's anycast address is delivered to only one of the nodes in the group, typically the node with the "nearest" interface in the group, according to current routing protocol metrics. This is in contrast with multicast services, which deliver packets to all members of the multicast group. Nodes in an anycast group are specially configured to recognize anycast addresses, which are drawn from the unicast address space.

Anycasting is a new service, and its applications have not been fully developed. Using anycast, an enterprise could forward packets to exactly one of the routers on its ISP's backbone. If all of a provider's routers have the same anycast address, traffic from the enterprise will have several redundant access points to the Internet. Should one of the backbone routers go down, the next nearest device automatically will receive the traffic.

Anycast has been proposed to allow end stations to efficiently access well-known services, mirrored databases, Web sites, and message servers. It can provide a versatile and cost-effective model for enabling application robustness and load balancing. For instance, anycast could provide enterprise robustness by assigning all the DNS servers in an enterprise the same anycast address.

4.14 Quality of Service

IPv4 carries a "differentiated services" byte and IPv6 carries an equivalent "traffic class" byte, intended for support of simple differentiated services. Both IPv4 and IPv6 can support the RSVP protocol for more complex quality of service implementations.

Additionally, the IPv6 packet format contains a new 20-bit traffic-flow identification field that will be of great value to vendors who implement quality-of-service (QoS) network functions. Such QoS products are still in the planning stage, but IPv6 lays the foundation so that a wide range of QoS functions (including bandwidth reservation and delay bounds) may be made available in an open and interoperable manner.

An additional benefit for QoS in IPv6 is that a flow label has been allocated within the IPv6 header that can be used to distinguish traffic flows for optimised routing. Furthermore, the flow label can be used to identify flows even when the payload is encrypted (i.e., the port numbers are hidden).

5. The Transition to IPv6

The transition from IPv4 to IPv6 could take one of several paths. Some are lobbying for rapid adoption of IPv6 as soon as possible. Others prefer to defer IPv6 deployment until the IPv4 address space is exhausted, or until other issues leave no other choice. Either way, given the millions of existing IPv4 network nodes, IPv4 and IPv6 will coexist for an extended period of time.

Therefore, IETF protocol designers have gone to great lengths to ensure that hosts and routers can be upgraded to IPv6 in a graceful, incremental manner. The transition will prevent isolation of IPv4 nodes, and also prevent "fork-lift" upgrades for entire user Populations. Transition mechanisms have been engineered to allow network administrators flexibility in how and when they upgrade hosts and intermediate nodes. IPv6 can be deployed in hosts first, in routers first, or, alternatively, in a limited number of adjacent or remote hosts and routers. The nodes that are upgraded initially do not have to be collocated in the same local area network or campus.

Many upgraded hosts and routers will need to retain downward compatibility with IPv4 devices for an extended time period (possibly years or even indefinitely). It was also assumed that upgraded devices should have the option of retaining their IPv4 addresses. To accomplish these goals, IPv6 transition relies on several special functions that have been specified by the “ngtrans" working group of the IETF, including dual-stack hosts, routers, and tunnelling IPv6 via IPv4. A dual-stack host is a computer able to handle both IPv4 and IPv6 packets. Such a computer can deliver packetized data to a single application that has been equipped to ask for data from both addressing domains. This facilitates easy transition from IPv4 to IPv6 since the application can then still receive data from its current communications partners, without change in any way noticeable to the users.

5.1 IPv6 DNS

Domain Name Service (DNS) is something that administrators must consider before deploying IPv6 or dual-stack hosts. In response to this issue, IETF designers have defined "DNS Extensions to Support IP Version 6". This specification creates a new "AAAA" (quad A) DNS record type that will map domain names to an IPv6 address.

Domain name lookups (reverse lookups) based on 128-bit addresses also are defined. Once an IPv6-capable DNS is in place, dual-stack hosts can interact interchangeably with IPv6 nodes. If a dual-stack host queries DNS and receives back a 32-bit address, IPv4 is used; if a 128-bit address is received, then IPv6 is used. Where the DNS has not been upgraded to IPv6, hosts can resolve name-to-IPv6-address mappings through the use of manually configured local name tables.

IPv6 auto configuration and IPv6 DNS can be linked by using dynamic DNS updates, coupled with secure DNS. By these means, DNS servers can be securely and automatically updated whenever an IPv6 node acquires a new address, enabling an additional measure of convenience compared with renumbering in IPv4 today.

5.2 Application Modification for IPv6

Applications that do not directly access network functions (i.e. do not call a socket or DNS API and do not handle numeric IP addresses in any way) need no modifications to run in the dual-stack environment. Applications that use certain interface APIs to communicate with the network stack

will require updating before using IPv6. For example, applications that access DNS or use sockets must be enhanced with the capability to handle AAAA records and 128-bit addresses. Applications which are expected to run both IPv4 and IPv6, as well as using IPv6 security, quality of service, and other features, will need more extensive updating.

Adding such a dual-stack architecture to all the existing hosts is, in fact, a significant effort. This effort has to be balanced against the benefits of IPv6, and against the effort to renumber the existing hosts if the network deployment grows past the restrictions resulting from insufficient address space.

5.3 Routing in IPv6/IPv4 Networks

Routers running both IPv6 and IPv4 can be administered in much the same fashion that IPv4-only networks are currently administered. Multi-protocol extensions to BGP4 have been defined by the IETF; one of them carries IPv6 prefixes. The IPv6 extension has been used widely in the 6bone since early 1997. It has been implemented by all the major router vendors and by the well-known gated daemon, and is described in a Standard Track document. IPv6 versions of other popular routing protocols, such as Open Shortest Path First (OSPF) and Routing Information Protocol (RIP), are already running. Administrators may choose to keep the IPv6 topology logically separate from the IPv4 network, even though both run on the same physical infrastructure, allowing the two to be administered separately. Alternatively, it may be advantageous to align the two architectures by using the same domain boundaries, areas, and subnet organization. Both approaches have their advantages. A separate IPv6 architecture can be used to replace the inefficient IPv4 topologies burdening many of today's enterprises. An independent IPv6 architecture presents the opportunity to build a fresh, hierarchical network address plan that will facilitate connection to one or more ISPs. This simplifies renumbering, route aggregation (summarization), and other goals of a routing hierarchy.

Initially, many IPv6 hosts may have direct connectivity to each other only via IPv4 routers. Such hosts will exist in islands of IPv6 topology surrounded by an ocean of IPv4. Therefore, there are transition mechanisms that allow IPv6 hosts to communicate over intervening IPv4 networks. The essential technique of these mechanisms is IPv6 over IPv4 tunnelling, which carries IPv6 packets within IPv4 packets. Tunnelling allows early IPv6 implementations to take advantage of existing IPv4 infrastructure without any change to IPv4 components. A dual-stack router or host on the "edge" of the IPv6 topology simply inserts an IPv4 header in front of ("encapsulates") each IPv6 packet and sends it as native IPv4 traffic through existing links. IPv4 routers forward this traffic without knowledge that IPv6 is involved. On the other side of the tunnel, another dual-stack router or host "decapsulates" (removes the extra IP header from) the IPv6 packet and routes it to the ultimate destination using standard IPv6.

To accommodate different administrative needs, IPv6 transition mechanisms include two types of tunnelling: automatic and configured. To build configured tunnels, administrators manually define IPv6-to- IPv4 address mappings at tunnel endpoints. Outside of the tunnel, traffic is forwarded with full 128-bit addresses. At the tunnel entry point, a manually configured router table entry dictates which IPv4 address is used to traverse the tunnel. This requires a certain amount of manual administration at the tunnel endpoints, but traffic is routed through the IPv4 topology dynamically, without the knowledge of IPv4 routers. The 128-bit addresses do not have to align with 32-bit addresses in any way.

Mbone deployment using IP-within-IP tunnelling has been quite successful, and validates this design approach as well as supporting the likelihood of smooth transition.

5.4 The Dual-Stack Transition Method

Initial users of IPv6 machines will require continued interaction with existing IPv4 nodes. This is accomplished with the dual-stack IPv4/IPv6 approach. Many hosts and routers in today's multivendor multiplatform networking environment already support multiple network stacks. For instance, the majority of routers in enterprise networks are multiprotocol routers. Many workstations run some combination of IPv4, IPX, AppleTalk, NetBIOS, SNA, DECnet, or other protocols.

The inclusion of one additional protocol (IPv6) on an end station or router is a well-understood problem. When running a dual IPv4/IPv6 stack, a host has access to both IPv4 and IPv6 resources. Routers running both protocols can forward traffic for both IPv4 and IPv6 end nodes.

Dual-stack machines can use totally independent IPv4 and IPv6 addresses, or they can be configured with an IPv6 address that is IPv4-compatible. Dual-stack nodes can use conventional IPv4 auto configuration services (DHCP) to obtain their IPv4 addresses. IPv6 addresses can be manually configured in the 128-bit local host tables, or preferably obtained via IPv6 auto configuration mechanisms. Major servers will run in dual-stack mode until all active nodes are converted to IPv6.

5.5 Automatic Tunnelling

Automatic tunnels use "IPv4-compatible" addresses, which are hybrid IPv4/IPv6 addresses. A compatible address is created by adding leading zeros to a 32-bit IPv4 address to pad it out to 128 bits. When traffic is forwarded with a compatible address, the device at the tunnel entry point can automatically address encapsulated traffic by simply converting the IPv4-compatible 128-bit address to a 32-bit IPv4 address. On the other side of the tunnel, the IPv4 header is removed to reveal the original IPv6 address. Automatic tunnelling allows IPv6 hosts to dynamically exploit IPv4 networks, but it does require the use of IPv4-compatible addresses, which do not bring the benefits of the 128-bit address space.

IPv6 nodes using IPv4-compatible addresses cannot take advantage of the extended address space, but they can exploit the other IPv6 enhancements, including flow labels, authentication, encryption, multicast, and anycast. Once a node is migrated to IPv6 with IPv4 compatibility, the door is open for a fairly painless move to the full IPv6 address space. IPv4-compatible addressing means that administrators can add IPv6 nodes while initially preserving their basic address and subnet architecture. Automatic tunnels are available when needed, but they may not be necessary when major backbone routers are upgraded to include the IPv6 stack. Upgrades can be achieved quickly and efficiently when backbone routers support full remote configuration and upgrade capabilities.

6. Myths

Because of its potential for future dominance and the number of detailed technical choices that had to be made, the birth of IPv6 has been attended by some controversy, and by a number of somewhat misleading stories that can distract network owners who are in the process of crafting their forward-looking network strategy. Confusion is to be expected, considering the implications of migrating our global internetwork infrastructure to an updated protocol. However, if the IPv6 myths are perpetuated indefinitely, there’s a risk that the Internet will not be able to progress beyond a patched-up version of IPv4. In these appendices, we try to counteract some of these myths.

6.1 Driving force behind IPv6 is address space depletion.

Many of the discussions about a new Internet protocol focus on the fact that we will sooner or later run out of globally unique network layer addresses, due to IPv4's fixed 32-bit address space. The various address registries that assign blocks of IP addresses to large network service providers and network operators have become quite cautious about the way these addresses are handed out, though most predictions for IPv4 address exhaustion target a time frame that starts well into the next decade. With the long haul in mind, IPv6 has been outfitted with a 128-bit address space that should guarantee globally unique addresses for every conceivable variety of network devices for the foreseeable future (i.e., decades). IPv6 has 16 byte addresses, or 340,282,366,920,938,463,463,374,607,431,768,211,456 addresses (over a third of a duodecillion of them, in fact). The number of addresses gets a lot of attention but it is only one of many important issues that IPv6 designers have tackled. Other IPv6 capabilities have been developed in direct response to current business requirements for more scalable network architectures, mandatory security and data integrity, extended quality-of-service (QoS), auto configuration, and more efficient network route aggregation at the global backbone level. These features are all specified with IPv6 in a way that would be difficult to realize as effectively in IPv4.

6.2 Can extensions to IPv4 replicate IPv6 functionality?

There have been multiple efforts to extend the life of IPv4 incrementally with evolutionary changes to the protocol standards and various proprietary techniques. One such example is the development of network address translators (NAT) that preserve IPv4 address space by intercepting traffic and converting private intra-enterprise addresses into one or a few globally unique Internet addresses. Other examples include the various QoS and security enhancements to IPv4, which are in general scaled-back or identical to mechanisms specified in IPv6.

We do not know how long IPv4's life can be extended by these techniques. What is certain is that the widespread introduction of NAT devices negatively affects the end-to-end viability of emerging Internet applications; in practice only; a limited set of well-known applications can be correctly handled by NAT devices or by application level gateways associated with them. In particular, NAT devices prevent the deployment of end-to-end IPv4 security. Furthermore, the development of new and innovative Internet applications is burdened with the design constraints posed by NATs. Since NAT is strictly unnecessary for IPv6, standard end-to-end IPv6 security can be deployed, and a future enlivened by new lightweight and more fully functional applications can be envisioned. NAT translation is also known to create great difficulty in the construction of Virtual Private Networks (VPNs), since it makes address space administration difficult and interferes with standard security mechanisms.

NAT also only works in a "flat universe" for a site accessing the global Internet - even moderately sized enterprises are not flat internally, with nested multi-party relationships. Realistic NAT deployment solutions would have to include routing via multiple ingress/egress NATs for load balancing, multi-NAT-hop routes and so on - all this would create in miniature the v4 (or in fact v6) architecture, since it is solving the same problem, but piecewise and badly.

It is hard to compare the costs of converting to IPv6 with those of remaining with IPv4 and its upgrades. Every network manager will have to make this comparison; but staying with IPv4 has been likened to the situation of a lobster in a pot of water, as the temperature slowly increases - at first, it feels comfortable.

6.3 IPv6 support for a large diversity of network devices

Over the next few years, conventional computers on the Internet will be joined by a myriad of new devices, including palmtop personal data assistants (PDA), hybrid mobile phone technology with data processing capabilities, and smart set-top boxes with integrated Web browsers, and

embedded network components in equipment ranging from office copy machines to kitchen appliances. Some of the new devices requiring IP addresses and connectivity will be consumer-oriented, but many will become integral to the information management functions of corporations and institutions of all sizes. These new devices require features not fully understood by most protocol designers during the initial growth of the IPv4 Internet.

IPv6's 128-bit address space will allow businesses to deploy a huge array of new desktop, mobile, and embedded network devices in a cost-effective, manageable way. Further, IPv6's auto configuration features will make it feasible for large numbers of devices to attach dynamically to the network, without incurring unsupportable costs for the administration for an ever-increasing number of adds, moves, and changes.

The business requirement for IPv6 will be driven by end-user applications. Applications for mobile nodes, electronic commerce, and those needing specialized routing features will be easier to design and implement using IPv6, especially as compared to IPv4 patched by NAT. To remain competitive in the coming era of high density networking, businesses should exploit IPv6 to create a highly scalable address space and robust auto configuration services that will remain viable in the face of an explosion of end-user networking needs.

6.4 IPv6 not primarily relevant to backbone routers

It is true that IPv6 address aggregation allows efficient multi-tiered routing hierarchies that prevent the uncontrolled growth of backbone router tables. However, many of the advanced features of IPv6 also bring direct benefits to end-user applications at the workgroup and departmental levels. For instance, applications will have available the mandatory IPv6 encryption and authentication services as an integral part of the IP stack. For mobile business users and changing organizations, IPv6 auto configuration will allow the efficient assignment of IP addresses without the delays and cost associated with manual address administration or even traditional DHCP, which takes place in many current IP networks. IPv6 is very much both an end-user concern and a business concern. This concern will become increasingly important as QoS flows and QoS routing become important architectural components of the Internet.

6.5 Asynchronous Transfer Mode (ATM) cell switching and IPv6.

ATM and other switching methods offer interesting technology for present and future internetworks, but ATM is, by itself, not a replacement for packet routing Internet architecture. ATM is better understood as a link-layer technology over a non-broadcast multiple access (NBMA) medium. It gives some isolation properties, and offers the promise for offering improved Quality of Service (QoS) connections for applications that need it. Even these hypothetical advantages are not yet fully developed for ATM, and it is possible that these advantages will be equally well available in future IPv6 networks not running over ATM.

Fortunately, network owners do not have to make a choice between ATM and IPv6 because the two protocols will continue to serve different and complementary roles in corporate networking. Large networks will make use of both protocols. For many network designers, ATM is a useful transmission medium for high-speed IPv6 backbone networks. Standards and development work is being devoted to integrating ATM and IPv6 environments. IPv6, like its predecessor IPv4, provides network layer services over all major link types, including ATM, Ethernet, Token Ring, ISDN, Frame Relay, and T1.

6.6 Use of IPv6 beyond Telecommunication Services

Some Internet pundits have characterized IPv6 as a concern that's outside the corporate network and outside the current time frame.

In reality, IPv6 is a standards track and mainstream solution for the operation and continued efficiency of day-to-day business activities. But the only way that IPv6 will take hold and succeed is if businesses and institutions of all types come to terms with the inadequacies of IPv4 and begin to lay plans for migration. In the past few years, Internet protocols have enabled a whole new style of distributed commerce that brings people together inside enterprises and gives enterprises access to the entire world. In fact, the sustained and impressive growth of the Internet, which has inspired the current engineering efforts for IPv6, is in large measure due to the penetration of the World Wide Web to business and consumer end users. Offering services to such end users is of interest to many more institutions than merely governments and telephone companies.

6.7 IPv6 in use with operating systems, applications, and programming techniques.

IPv6 obviously requires certain modifications to the network protocol handling modules installed on the relevant computers. However, this typically requires little or no change to the base operating system. Simple and natural modifications, typically confined to fewer than a dozen lines of the programs, can be made to enable applications to use IPv6 addresses directly. Since IPv6 reserves a part of its address space for compatibility with IPv4 addresses, applications modified to handle IPv6 addresses can still communicate with existing IPv4 clients and servers.

Moreover, the transition strategies defined for IPv6 deployment within the IPv4 Internet should make the gradual adoption of IPv6 a smooth process that allows existing applications to be converted for native IPv6 operation in a gradual, controlled manner.

6.8 IPv6 Benefits

IPv6 appears as an incremental enhancement to IPv4, and some people say that if we are going to go to all the trouble to switch network-layer protocols, we really ought to go all out for some really futuristic feature-full new protocol. This argument ignores the following simple facts:

- The purpose of a network-layer protocol is to hook together networks, and IPv6 builds on the amazing success of IP, by not forgetting the successful parts, and by repairing the known faults. This is far different than starting over again with something unknown and untested.
- Those who claim that it is too early for IPv6 ignore the facts that existing solutions extending the life of IPv4 are clearly stopgap measures, and that one can put IPv6 into service now.

6.9 Renumbering in IPv6

Although IPv6 has gone a long way to enable more convenient renumbering operations, it is a mistake to say that renumbering is a completely solved problem. IPv6 engineers are still considering designs for renumbering routers, and for renumbering collections of computers larger than a single network. Furthermore, applications that have been ported from IPv4 to IPv6 do not automatically become more able to support renumbering. Some applications will require small design improvements in order to support renumbering. Lastly, the biggest impediment to renumbering seems typically to be the institution of administrative practices that key information directly on IP addresses instead of some more appropriate indexing method. These administrative practices require attention and adherence to more modern guidelines for Internet administration before the problem of renumbering can be considered to be solved.

6.10 Routing in IPv6

IPv6 offers improvements for routing in a number of ways. It allows for allocation of IPv6 addresses in a way that is more favourable for aggregation than existing IPv4 allocations. It allows for more streamlined packet forwarding than IPv4 routers can do, especially when IP options are used. IPv6's larger address space offers opportunities for more optimal network planning, since the constraints for planning out network connectivity have been relaxed to such a great extent. Furthermore, since every IPv6 router can be presumed to have security processing enabled, it is much easier to institute the appropriate security measures for authentication and keeping private data private.

However, there are still many operational issues that need attention.

- IPv6 routing protocols are largely adapted from almost identical, IPv4 routing protocols, and thus inherit some of the same problems.
- Improvements continue to be made to routing protocols to improve their stability, convergence time, and configurability. One of the hardest problems is to make routing protocols more human-friendly, so that it does not take a genius to make the routing fabric work reliably. There are remaining issues surrounding multi-homing that have not been solved. All of these issues will continue to receive the attention of engineers involved with the creation of IPv6. The scoped addresses and native security are expected to make their solution much easier.

7. New IPv6 Policy Framework

Recently new developments concerning IPv6 address allocation policies have been started within the LIRs. This proposal is based on requirements expressed by the RIPE community as well as those of the other RIR communities. It provides a set of proposed policies for the management of IPv6 address space, specifically concerning the allocation of address space allocated by Regional Internet Registries (RIRs) to organisations operating IPv6 networks. Under the current system of management of global IP address space, Regional Internet Registries (RIRs) are responsible for allocation of address space to organisations within their respective geographic regions.

In 1999, the RIRs APNIC, ARIN and RIPE NCC published a provisional policy document for IPv6, which has been in operation since then. Since 2000, this document has been under review and discussion, and through this process, many issues have been raised. It is the aim of this document to propose a new policy framework for IPv6 address space management which takes into account the operational experience of the past 3 years, and addresses most if not all of the major issues raised through the open review process

This framework, specifically takes the following into account:

- The allocation criteria should be such that it is easy to obtain IPv6 address space.
- The size of the initial allocation should be large enough to allow flexibility in addressing infrastructure and customer sites.

This results in the following changes to previously discussed IPv6 allocation criteria:

1. to recognise existing infrastructure (both IPv4 and IPv6) where it exists and calculate IPv6 address needs based on existing networks.

2. to apply the slow start mechanism only for 'IPv6 only' networks without existing IPv4 infrastructure

3. to reduce the minimum allocation size for those IPv6 only networks (unless larger requirements are shown)

The above is not to attempt to provide details of policy implementation, procedures or documentation; nor does it document requirements for management of address space, which is allocated. These policies could be established globally or regionally as appropriate, based on global consensus regarding the fundamental principles described here.

7.1 Address Space Requirement for Initial Allocation

It is proposed to recognise existing network infrastructure and address utilisation (both IPv4 and IPv6). New IPv6 address needs are then based on these existing networks.

In assessing a request for an initial allocation, there are 3 possible cases to consider:

- the requesting organisation has an existing IPv4 network which will be addressed by the new IPv6 allocation
- the requesting organisation has an existing IPv6 network
- the requesting organisation has no network at all

7.2 ETNO Common Position to IPv6

ETNO is the principal European trade association of network operators, gathering 45 operators and LIRs at the same time (www.etno.be). It includes both historic operators and new entrants. Some of them already have been granted IPv6 addresses.

The following position has been signed up by all 45 members of ETNO, further to the drafting work led by Niall. That is available on the ETNO site at http://www.etno.be/ETNO positions/latest position papers. ETNO supports /29 because it is consistent with the current allocation policy as defined by the RIRs and better fits the LIRs' needs than the current /35. Indeed, ETNO considers /29 as a minimum allocation size to allow a sub-TLA registry to run its activities efficiently.

The rationale for a wider address block for operators acting as LIRs lies in aggregation, efficiency of routing tables and better management of address space by TLA registries. These registries need enough space to serve their customers and NLA registries in the best possible way.

The foreseeable needs for public addresses for the coming years, especially for UMTS, DSL, cable modems... is a source of concern for operators and the Internet community as a whole. It can be expected that the IPv4 space will reach exhaustion in the most likely case around 2005 - 2007. Therefore, it is in the Internet community's interest to stimulate at the soonest a broad usage of IPv6 addresses, where no shortage should exist in theory.

It is the ETNO opinion that such stimulation can only be obtained by setting up rules of management of IPv6 addresses that would avoid at the same time:

- to introduce artificial limitations that might hamper the move to IPv6, and exaggerated increase of routing tables.

ETNO favours a structure that permits maximum aggregation. That would enable operators and ISPs to structure their addressing plan in a long-term perspective.

ETNO strongly supports the allocation of /29 to all TLA registries, including the existing ones.

In the same spirit of facilitating a timely move to IPv6, ETNO also suggests that the opportunity be taken at the end of the bootstrap period, expected end 2001, to review and ease the conditions for registries to be allocated IPv6 addresses.

8. Numbers & Numbering and Names

The International Telecommunication Union (ITU) coordinates Internal Numbering. The ITU provides a means to plan the development of the global telephone numbering space It has been defined the format of numbers in the ITU-T recommendation E. 164 and assigns country codes. The transition from IPv4 to IPv6 requires careful planning on the parts of the users and public authorities alike. It is therefore, important that whatever, proposals are derived out of the work of the IPv6TF working Groups be communicated with the ITU-T, ENUM, INTUG, 3GPP etc.

9. Recommendations to Member States

According to the political mandate given to the IPv6 Task Force to elaborate recommendations of European policy towards implementation of IPv6 in all industry sectors in Europe; the Mobile Wireless WG sees the following recommendations where IPv6 is indispensable for the *future networked economy* in all European member States:

A.1 It should be understood that the move towards native IPv6 is a major step for Europe to gain back its dominant position in the Mobile Industry. With the Asian industry already investing millions of US dollars in IPv6 enabled consumer appliances, Wireless devices and in other industry sectors, Europe needs to do the same.

A.2 Moreover one needs to stress that IPv6 is not limited to fixed core networks, wireless and cellular systems. IPv6 deployment will take place in different Industry sectors creating new applications and services to be networked and could possibly be delivered by 3G networks. IP exist today on top of many technologies and with IPv6, it is expected the IPv6 with carry everything on top. Europeans should be prepared to invest resources and implement their strategies towards a smooth medium to long-term transition.

A.3 European Union and Member States should raise awareness of IPv6 within appropriate organisations and business entities, and especially SME companies. Education and public awareness among all relevant parties is seen as the key success factor in order to facilitate successful early transition towards IPv6.

A.4 The European Union and Member States should not only support the wide spread use of Internet across Europe but also encourage the integration of IPv6 in other next generation technologies through appropriate mechanisms. All artificial or strict regulations, time-lines, fees or regulatory mechanisms should be avoided. Hence, allowing for Network operators and ISP's to offer more Cellular services to their users.

A.5 On global bases, IPv6 could play a vital role of globalisation. It is therefore, also important that the European Member States in their messages across to the Developing countries express the needs of harmonization and the economical benefits of using a wide spread technology in both the telecommunications and Internet sectors.

10. Recommendations to the Commission

The European Commission could play an instrumental role in coordinating some of administrative Issues for Regulators, Policy Makers and Numbering Authorities if the integration of IPv6 in the Internet is not to be unduly constrained. However, the management of that integration needs to handle with care and sensitivity if it is not to be disruptive or cause unnecessary costs for users. Therefore,

B.1 The European Commission should initiate a discussion of IPv6 policy matters at a European level. The Commission should assist and encourage the wide spread deployment of IPv6 within the EU member states in a timely Manner.

B.2 The European commission should assist in assuring that IPv6 Addresses are made available to all sectors of the industry including UMTS operators in the numbers needed and at reasonable prices. In order to do so, an agreement between the Commission and RIPE should negotiate as to how this can take place.

B.3 The European commission should provide guidelines along the recommendations of the Mobile interconnect directives to RIPE and ISP's.

B.4 The European commission should assist and encourage industry (terminal, Infrastructure suppliers, application platform and Software vendors) to accelerate their IPv6 product roadmaps and the development of IPv6 enabled solutions, in order to enable full deployment of IPv6 networks as soon as possible.

B.5 In the Research and Development Programme, the European Commission should put a high emphasis on the use of IPv6 and integration of IPv6 Products in the Projects.

B.6 It is recommend that the European Commission:

- Support and encourage market studies that will provide basic information so that the business case and timing for the integration of IPv6 can be determined;
- Assist and encourage the development of an overall technical strategy which is supported by Vendors and Operators, and leads to the introduction on the basis of sound commercial criteria and product availability

B.7 The 3GPP specifications draw heavily upon work in other bodies (e.g. IETF). It is therefore recommended that the 3GPP standards and the Draft RFCs of the IEFT be aligned in order to meet the needs of Operators to deploy the IMS, and both "traditional Telco" and "more IETF oriented IP technology" Vendors to take advantage from this new business opportunity. The European Commission should therefore, initiate and encourage larger European activities in all relevant forums and Research communities and workgroups, particularly in ITU to harmonize the requirements to be brought forward to the standards bodies.

B.8 The European Commission should develop joint activities with the ITU Secretariat, particularly the ITU Telecommunications Development Bureau, in order to raise awareness among the ITU Member States of the importance of IPv6, particular for 3G mobile communication systems, discuss issues related to the integration of IPv6 etc. The ITU could in return assist in supporting a global strategic action plan.

B.9 The European Commission should encourage its Member States to contribute input documents to the relevant ITU-T groups on the use and importance of IPv6.

11. General Recommendations to the Industry at large

These recommendations are seen to address the Industry at large and all the involved parties (RIPE, ETSI, ITU, 3GPP, UMTS Forum, GSMA, ISPs associations including ISPs etc.)

C.1 With a view to implement IP Multimedia under Release 5 in 2003~2004, industry will be requested to submit contributions to 3GPP to accelerate the pace of development of specifications work on IPv6 for 3G mobile communication systems (UMTS). This is an industry-wide issue where since two years the UMTSF has been active in close relationship with the IPv6F. It is now time for this to be taken forward and the Task Force

make specific recommendations on how best to progress the introduction of IPv6, the requirement for address allocation, and the steps needed to obtain it.

C.2 Release 5 is currently scheduled for completion by the 3GPP in March 2002 with subsequent functionality releases in October 2002, Whilst a recommendation to accelerate this work is fully supported, it is likely to have limited impact upon the current 3GPP Release 5 activity. The 3GPP specifications draw heavily upon work in other bodies (e.g. IETF) – it is therefore recommended that the Task Force study the capabilities of the 3GPP standards and the Draft RFCs of the IEFT meet the overall scope of the Operators and deployment of IMS. Any deviations should than be deferred perhaps to future Releases (6, 7…).

C.3 It is well understood that due to 2.5G technology availability, Operators were forced to make significant investments in legacy IPv4 infrastructure but also have deployed this in other non-cellular areas. Here again, IPv4 deficiencies being availability and cost of addresses for some 550 Cellular operators around the world. However, these networks will require further upgrades and deployment as the market grows, the operators should consider the network update with IPv6 for easy of transition to IMS. The backward compatibility / transition with coexistence of both IPv4 and IPv6 and no particular preference has been in the mean time well defined by a number of vendors. The Task Force can attempt based on some scenarios to provide guidelines.

C.4 With the introduction of IMS around 2004~2005, 3G Operators will have to have IPv6 deployed in their networks at the latest. The Task Force should recommend a technical strategy fully supported by vendors to execute this on the basis of sound commercial criteria, commercial factors, product availability and the significant investment made by operators and service providers. The Task Force will produce the necessary requirements to study the market and provide the basic information for each industry body to individually derive its business case and time for transition to IPv6.

C.5 Network operators will be the first ones to benefit from an early competitive edge when IPv6 is rolled out, and their business models affected by the immense opportunities that IPv6 would offer. Nevertheless, Network operators should be allowed in the light of their business interests, to make the decision as to when they will introduce IPv6 in their networks. This should also take into account the legacy users, systems and networks.

C.6 Industry should consider in their manufacturing plans that the majority of mobile devices will require some form of IP connectivity. The simplest way to offer these mobile devices the fullest range of services is to have a unique globally routable IP address available for each terminal; the same applies for all addressable Network components.

C.7 The Task Force will continue to complete its mandate as described in its ToR and produce the detailed output documents of

1. Requirements for a market study
2. Base line document for naming and addressing
3. Roadmap to align standardisation activities (IEFT, 3GPP, UMTSF etc.)

EU IPv6 Task Force discussions with industry, operators and the research community have focused mainly upon the above issues.

Is a general, one should note that IPv6 technology is already available on the market today. Moreover, gradually, the industry is beginning to understand that there are several enhancements in IPv6 over the currently available IPv4. Trial projects and on-hands experience are on going but what is yet to be accomplished is a mass educational program.

APPENDIX A: POSITION PAPERS ON IPv6 AVAILABLE FROM COMPANIES
(Non-comprehensive list) Also, see the WG 1 Report.

http://www.ipv6.org/draft-iab-case-for-ipv6-06.txt

References:

The Case for IPv6 draft-ietf-iab-case-for-ipv6-06.txt