

Arden L. Bement, Jr., Director
National Institute of Standards and Technology

Michael D. Gallagher
Acting Assistant Secretary for Communications and Information
National Telecommunications and Information Administration

In response to:
Request for Comments on Deployment of Internet Protocol, Version 6

I appreciate the opportunity to comment on this issue of importance to the public good. This is a personal response. While I have several roles related to IPv6 development and deployment, these personal comments may or may not appear in whole or in part in any responses from those organizations. Specifically, this response *MUST NOT* be interpreted as an official position from Cisco Systems, the North American task force of the IPv6 Forum the IETF, or for any other group where I may have provided IPv6 guidance or instruction.

As an active participant in the IETF, I have been involved in IPv6 development, implementation, and operations since its inception. My varied personal roles include operational management at a Federal R&E network; IPv6 program manager for a leading host operating system stack; IPv6 technology leader at Cisco Systems; co-chair of the IETF transition tools working group; technology director of the NAv6tf of the IPv6 Forum, as well as past member of the Internet Architecture Board.

While it is clear that substantial effort went into the preparation of detailed questions, I will provide summary answers to the four main questions. One reason is due to the answers for several questions are situation dependent, when the question appears to request a 'one-size-fits-all' answer to very different problems. Another reason is the apparent misalignment between the tone of the questions with likely deployment and transition strategies. Finally repeating context in the answers to the individual questions seemed redundant.

In response to one specific question; "...will late entry into global IPv6 markets by U.S. firms have a significant long-term negative effect on market shares and economic performance?" Japan, China, and to a limited extent, Europe, already see the lethargy in the U.S. as an opportunity to seize economic leadership. The applications being developed there will not only be cheaper than the complex IPv4/NAT-based ones in the U.S., the developers will have more creative freedom as they are free to focus core products and services, or those items people will pay for, instead of being forced to spend time on context, designing workarounds in the limited IPv4/NAT environment. Application leadership is a natural outcome awarded to those with the most efficient development process, and the lowest ongoing operational cost. In the global realm of competition, once leadership in Internet services and applications is lost, it will be difficult for U.S. application developers to regain.

Respectfully,
Tony Hain

tony@tndh.net
<http://www.tndh.net/~tony>

(1) The benefits and possible uses of IPv6

Larger address space

The primary benefits of IPv6 are derivatives from is the larger address space. As will be discussed later, the larger address space afforded by IPv6 enables inclusion and participation by all. There were a few design changes to streamline packet handling, but the significant change was to increase the address size. One less than obvious benefit of the larger address space is the simplicity it brings to plug-n-play configuration for the emerging vast number of consumer appliances¹.

For the enterprise the larger address space of IPv6 simplifies the process of subnet management. One of the hidden costs within corporate networks is the time and effort needed to change the address space allocated between subnets. Growth or shifts in function location will frequently cause a mismatch between the number of active nodes and the amount of address space allocated to a specific office. The typical response is to have several people spend a weekend manually reconfiguring the impacted office. This problem does not exist with IPv6 because every subnet will have substantially more address space allocated than the number of nodes that any media will support.

The larger address space allows removing NAT from the network, which in turn allows application developers to have more creative freedom. With this freedom they are allowed to focus core products and services, rather than network logistics. As noted in the questions, the original Internet assumption was open peer-to-peer. Almost all application developers still assume this to be the case, but reality shows that everyday the network is segmenting into additional independent addressing realms. There are multiple small armies of development teams working on the hard problem of how to traverse the multiple variants of NAT in a way that allows the application developer to still believe there is a single network. This wasted effort often needs to be reproduced for each new application. Even when NAT traversal schemes are marginally successful, if cooperation of the NAT is required, all devices along each given path need synchronized upgrades to enable application deployment.

NAT

The intensive list of questions about NAT exposes the overall confusion created by product marketing campaigns. For a discussion of the architectural impacts please see IETF RFC 2993. The most important issue to point out is that as a technology, *NAT provides absolutely no security or protection*. The process of mangling headers does absolutely nothing to stop the attacker, and actually simplifies some attacks by providing a single point of focus. The ‘feature’ the marketing campaigns confuse in NAT products is really a router that only forwards after state is established. Stateful routers existed long before NAT was published as an IETF document. Moving the market onto an IPv6/stateful-router requires the marketing organizations that are currently pushing NAT to recognize the difference, and then promote an IPv6/stateful-router as a better solution to hacking, that doesn’t in turn break the cool new applications.

To the issue about NATs precluding peer-to-peer devices and applications; as the references suggest, network researchers can develop such technologies, and there are a few NAT traversal technologies in the market today. Some include middle boxes acting as relays (which raise performance and scaling concerns), while others use tunneling techniques to effectively

¹ Mario Tokoro, Sony’s Co-CTO discussed Sony’s strategic product direction.
<http://www.ipv6style.jp/en/interviews/20030212/index.shtml>

bypass the NAT. The difference between the ability to create these technologies and their viability in the market lies in the ongoing costs associated with operational complexity. For example, last week I was in a hotel using one of the tunneling mechanisms, where the virtual tunnel interface was allocated exactly the same addresses as the hotel was using. This duplication prevented any further network activity. When situations like this happen to the average network user, they will create an expensive call to a support desk, and even then the problem may not be resolvable. Either way the process results in lost productivity to the network user, as well as a need to increase support staff at the service provider. Another issue is that the operational complexity of synchronizing all the end points & NAT devices along an arbitrary path is a substantial and ongoing cost for the enterprise or consumer. When different applications require different changes, the complexity of managing the resulting matrix and making sure the complete set of work-arounds actually supports all of the applications the consumer cares about is very high and never ending.

When middle-boxes add real value, the market will adopt and sustain them. When they simply get in the way and break applications, these middle-boxes may find initial acceptance due to projected value, but over time as the operational problems become acute, their value diminishes. Given the limited pool of IPv4 addresses, NAT has real market value today by allowing the network to expand unconstrained by restrictive global allocation policy. Unfortunately this comes at a cost, by limiting the kinds of applications and services that are possible. IPv6 provides address expansion while enabling applications, so the IPv4/NAT technology is left with no marketable value.

Another impact of NAT is the accessibility of devices on the network. While NAT can be transparent to the client/server applications like web browsing & email, that is only the case when the server is located in the public address space. Case in point, both my son & I run independent web & email servers for our own uses. While we could configure a NAT to make -one- of those work, there is no way to share a single address & application port for both of our servers to operate at the same time. This exposes the tip of the iceberg as things like voice over IP bring telephony to the Internet appliance space, and the number of devices that will end up in the home needing to be reached over the same application protocol identifier exceeds the one available to the public side of the IPv4/NAT. Multi-player games, using a PDA remotely to access to personal files, and home appliance diagnostics are yet other instances.

The question about NAT that DoC should be paying attention to is not 'how can we develop work-arounds to delay the inevitable?', but 'how soon will we lose the lead in developing applications for e-commerce and the Internet?' When other countries find it is cheaper to build and operate IPv6 based applications, it won't matter that there are IPv4/NAT work-arounds available from companies in the U.S.

QoS

At this time, a specific non-benefit of IPv6 is quality of service (QoS). Currently, QoS in IPv6 is exactly the same as IPv4. There was an additional field added to the IPv6 header that is intended to assist with QoS, but its definition is still moving through the standardization process. Until that field is agreed on, and implementations start using it, the QoS capabilities will be the same for both versions.

Security

Another area that is over hyped is Security. In general, the IP layer is about movement of data between endpoints. As such, IPv6 does not by itself improve overall system security,

particularly with respect to IPv4. Fundamentally, 'security' is an ambiguous concept, and different people will answer these questions based on what they think is important to their wellbeing. From one perspective, when there are keys available the ability to easily use IPsec to protect all applications improves overall data integrity and confidentiality for the end user. From another perspective, enabling IPsec to work in more cases is considered to reduce security by those that believe they gain security from being able to monitor the content of each packet stream. This is really a matter of trust, and not a protocol version issue, because alternative encryption technologies like TLS & HTTPS also mask the data stream to the detriment of those who believe in content monitoring.

In either case, globally unique IPv6 addresses bring customer traceability back to the table, so in the area of assisting law enforcement, the DoC, along with other appropriate agencies, should work with the service providers toward a plan of static prefix allocation to each customer. The significantly larger address space of IPv6 allows prefix allocations to be static without concern for exhaustion. At the same time, agencies should know that there is a capability in IPv6 called 'Privacy addressing' that is intended to deter malicious or unwanted marketing web sites from tracking mobile appliances as they move around the network. These addresses will make the mobility aspects of end point tracing as difficult as they are with IPv4/NAT. While the prefix part of these addresses are traceable to the ISP/customer demarcation point, tracing beyond that requires access to any local routers. As a result, law enforcement should not plan on being able to trace a specific node as it moves between attachment points, or over extended periods of time.

IPsec

IPsec is only one component of an overall security model. The DoC asked; "Are there critical IPsec implementation issues that are independent of the version of IP employed?" The most critical deployment issue is the lack of a public key infrastructure (PKI). At this point, this really is a non-technical problem of coordinating trust relationships. Even if limited communities can deploy private keying using PKI technologies, for general use there will still be a need to tie the communities together. Since that is a global problem, governments will need to be involved.

New opportunities

One of the untapped potential new markets is in the area of remote service & support. One problem that the fractured NAT address structure creates is the inability to unambiguously reach the Internet enabled appliance from the remote service center. With the global uniqueness of IPv6 addresses, it will be possible for appliance manufacturers to remotely monitor and diagnose problems. This has the double benefit of only requiring a truck roll when there are real physical problems that require hands-on, and allows that truck to be carrying any expected replacement parts on the first trip. Deriving the savings to the appliance manufacturer is straight forward, but the substantial savings to the consumer in reduced hassle at having to arrange time off from work to meet the repair truck are incalculable.

A new characteristic of IPv6 that is just not possible with IPv4 is that multiple parties can maintain separate simultaneous access paths without real-time manual intervention. One example would be the manufacturer of transportation equipment can maintain an address range for the maintenance uses suggested above for appliances. At the same time, the company operating that equipment would have its own address range for operational management. Yet another block could be assigned to any regulatory body, while still another could be dynamically allocated to the vehicle by the current port of call. Each of these address blocks could be implemented with different routing and access policies, and this is significantly beyond anything possible with the limited address pool of IPv4.

Additional 'new services' might be; the ability to deploy servers in the home like a personal web site for sharing vacation or baby pictures among the family, or the ability to simultaneously reach into more than one home system to retrieve a file (using a PDA on WiFi at Starbucks, Joe realizes he needs a document from his home machine, at the same time Joe's daughter is visiting a friend's house and needs access to homework on her machine), the ability to conduct multi-player games without the need for, or latency unfairness of a central server (geographically diverse family members share a 'game-night' experience as if they were around a table). While all of those services could be accomplished using IPv4 without NAT, the reality of a limited address pool is forcing wider deployment of NAT. This in turn removes the possibility of not only these applications and services, but any other new ideas where address uniqueness is required. In short, an IPv4/NAT Internet is restricted to the types of services where clients around the edge connect to servers in the middle; while IPv6 allows for any node to directly interact with any other node, so there are no limitations on the field of potential applications or services.

3degrees² when coupled with the IPv6 stack in Windows XP is an example of an application that works by using IPv6, independently from any IPv6 service offering. The application is only aware of IPv6, and thereby sees a consistent addressing environment, while the underlying operating system stack takes care of using an IPv6 service if it exists, or tunneling over IPv4 when service is not available. Basically any application could use this same approach, so a partial 'deployment' of IPv6 does not represent a problem. A partial 'implementation' (as in code in a specific device) might or might not allow an application to work at all, much less thrive.

It should be noted that only the network manager and application developer will know the difference between the bit patterns which identify IP versions. Consumers will see application cost, operational complexity, and performance as the only differentiators. They specifically should never be aware of the IP version; just that they can get their task done in the most economical and expedient manner.

(2) Current domestic and international conditions regarding the deployment of IPv6;

State of IPv4

As of 2/10/04, the ARIN report³ to NANOG showed the IANA registry central pool with 89 remaining reserved /8 blocks, or 34.7% of the total. (/8 is the typical allocation unit to the regional registries⁴). You will undoubtedly receive many pointers to Geoff Huston's work⁵ analyzing IPv4 address consumption. While his work documents current trends, the measurements are taken over a relatively short time window, during a period of reasonably strict allocation policies. Due to the current RIR allocation policies, we are in a period of managed growth. Geoff's numbers demonstrate that we can manage the remaining lifetime of the IPv4 pool through the strictness of the allocation policies, thus with even more restrictive policies we can extend the lifetime of IPv4 indefinitely. The outcome of all of this is that it is unlikely we will ever allocate the last IPv4 address, but from the perspective of consumers being able to accomplish their goals, we are effectively already out of IPv4 addresses. People are deploying NAT as a means to acquire the address space they want, but can't otherwise economically obtain.

² <http://www.3degrees.com>

³ <http://www.nanog.org/mtg-0402/plzak.html>

⁴ <http://www.iana.org/assignments/ipv4-address-space>

⁵ <http://www.potaroo.net/presentations/2003-09-04-V4-AddressLifetime.pdf>

In the long run the policies will simply become so restrictive that only the very privileged few will have access to public addresses and the kinds of applications that require those.

When one considers achieving mass-market status where there are claims this requires reaching 20% of the target market, it is clear that IPv4 will be incapable of achieving that goal. Since only 36 of the 208 (17.3%) UN recognized entities have reached at least 20% of their populations, on a global scale IPv4 hasn't even achieved the defacto mass-market goal. Bringing just the top 15 countries up to the point of a single address for 20% of their population would require over 400% of the remaining IPv4 pool ⁶.

Another area of consideration is that much of the historical allocation of IPv4 space was during the transition from multiple people sharing an IP address on a central main-frame, to widespread use of personal computers with individual addresses. In the consumer space these addresses were typically time-shared via dial services, which are now also migrating toward a one-address-per-person usage as people migrate to the always-on broadband technologies. Then as we look forward, we already see early adopter homes with multiple Internet enabled appliances per person. Where in the past there might have been a single home computer with directly attached devices like a printer, now there are emerging cases of multiple computers in the home controlling an array of network attached appliances. Given this trend, even by restricting Internet service to those that already have it, we further deplete address space just dealing with the shift in address consumption per person.

While multiple IPv4 technologies including NAT, CIDR, and DHCP slowed the consumption of address space, IPv4/NAT is a double edged sword ⁷. On one hand NAT allows for expansion of the number of attached devices, at the same time it creates a problem for application developers. As noted in the questions, there are concerns about NAT as a single point of failure. While in the typical case NAT would represent a single point of failure, in the vast majority of those cases there is a single transmission path anyway, so the NAT by itself didn't really create the single point problem. In the relatively few cases where there are alternate paths, the technology exists (though imperfect) to move state between NATs to deal with circuit failures.

As will be noted in an upcoming Internet Draft, in recent months I have been approached by multiple U.S. firms that have used up the space allocated for private use behind the NAT, yet need additional space to continue growing. They would acquire public address space from the RIR pool, but have expressed concern that the cost to do so is prohibitive. As globally routable address space is a public resource, it needs to be economically available to all.

State of IPv6

It will not be necessary for government agencies to take the approach of defining "mandatory" and "optional" services within the IETF standards. Experience shows that there have not been any obvious problems that arise in implementing the IETF standards for IPv6, as major operating system and router vendors already have implemented and demonstrated interoperability at periodic events. Even the transition technologies have proven they do not create an undue hardship on equipment or software providers as they are frequently cheaper than the alternative of developing IPv4/NAT work-arounds.

⁶ http://www.nav6tf.org/RIR_eNations/Geo-Politics_IPv4_Gap.pdf

⁷ <http://www.ietf.org/rfc/rfc2993.txt?number=2993>

As products are procured in the market place, the initial metric to gauge IPv6 deployment that matters is the ability to use IPv6 enabled applications. Since the current volume of deployed native IPv6 and IPv4 network equipment is proprietary information, a precise count can't be known. Even if it could it is not likely the respondent could answer in a public commentary. If any appear, such measurements require a consistent context with respect to comparable measurements to the IPv4 market. During transition, as long as IPv6 traffic tunneled over IPv4 is counted as IPv6, and not double counted, traffic volume would be one viable measure of deployment. As the market matures, the efficiency of the overall system will become an issue, where the offering of native services might provide alternative measurements. In the short term, the simple ability to use applications like VoIP or peer-to-peer games between arbitrary endpoints is a sufficient gauge to define IPv6 penetration.

Domestic vs. International deployment

In general, North America lags the world in terms of deployment and focus. The common perception is that this is due to the abundant availability of the IPv4 resource from historical allocations. At the same time there is a strong NAT deployment in the U.S. by consumers as ISPs restrict availability (charging by the quantity of addresses in use) due to the real limitations of this supposedly abundant pool. Japan & China are putting emphasis on IPv6 deployments as a catalyst to spur economic growth. The Asian region has the most pressing need for addresses due to the disparity between the global population distribution and historical address allocations. Thus it is no surprise to see the initial commercial services emerging in that region, complementing the strong emphasis on application development ⁸.

Since commercial deployments have already started in parts of the world, it is clear that the current set of IETF IPv6 standards are technically complete enough to enable viable and widespread commercial deployment. **If there are inadequacies in the ability to deliver something specific to the U.S. market, those will only surface when the U.S. market actually deploys IPv6 in quantity for production use.**

Market segment

The major end host operating systems include IPv6. In the U.S. the routers that are deployed and IPv6 enabled are primarily in the R&E, and Fednet backbones. Typically the campus networks attached to those have not yet enabled IPv6 to the researcher. There are also a small number of IPv6 enabled routers in ISP environments. Identifying the approximate domestic and global value of all deployed IPv4 and IPv6 equipment is difficult because most IPv4 equipment is capable of IPv6 with a software upgrade. Only the very high-end equipment with 32 bit restricted hardware is truly IPv4-only, and the box count in that pool is so small that answering the question steps into proprietary protected data. A different approach to the question of market value would be; the useful value of all IPv4 equipment where NAT prevents completing the task at hand is zero, while enabling the task through its IPv6 counter part is priceless. The absolute dollar values make little sense without a context of the relative value of being able to accomplish a task or not.

During this period of limited focus, U.S. suppliers are meeting current domestic demand. At the same time, the segments of application development, commercial service, and training are behind the curve with respect to global demand, and will find it challenging to ramp up fast enough when the domestic demand takes off. Specifically, routers have been available from

⁸ <http://www.v6pc.jp/apc/en/invitation.html>

multiple suppliers for several years, new operating systems shipped since early 2001 have all had IPv6 available in some form (a few suppliers are still limiting distribution, but this might have more to do with applications and packaging than anything else), while the applications segment is clearly lagging and stands to lose their lead to Japan where applications are a major focus. Commercial service providers are doing their normal search for someone willing to pay, but even with customers at hand in the short term services will be limited by a lack of trained staff.

Despite the obvious effort put into developing this set of questions, the DoC was unclear in its point about IPv6-capable. Does 'capable' mean the software & hardware is 'available', or 'installed & turned on'? Virtually all routers and general purpose operating systems will handle IPv6 packets when current IPv6 capable software is installed and enabled. This is as much a training problem as it is a perception of need. In the U.S. the historical IPv4 allocations to commercial enterprises and Universities was very generous, so the perceived need in those markets is relatively low. Since consumers are already being charged per address, their perceived need is slightly higher. While the number of consumers that have IPv6 service available to them is close to zero, as will be discussed in the transition tunneling technologies, it is fortunate this doesn't really matter. Enterprise customers that persist in seeking IPv6 service from their ISPs are able to acquire it from multiple providers (only Verio has announced ⁹ commercial availability, but there have been private reports that others will provide it if the requestor is serious). If all customers started demanding IPv6 at once, a portion of the backbone routers might need out-of-cycle hardware or software updates to align their efficiency of IPv6 traffic handling with the demand. In general, router deployments will follow application availability.

The need for IPv6 in the U.S. is becoming acute as private conversations with wireless carriers shows that the restrictive IPv4 allocation policies are inadequate to meet their business plan growth rates. It is likely services from that market segment will emerge after their normal design/install/support-training cycle, around 2007. The DoD 6/2003 directive to acquire IPv6 capabilities has piqued the interest of their contractor community, so demand from the related segments will pick up over the next few years.

Enabling IPv6 can be as simple as installing an application. The 3degrees application is an example where the consumer acquires an application for its features, and it interacts with the operating system to enable IPv6 along with any necessary transition technologies. Unless the consumer is really interested and digs deeply into the support part of the download site, they will remain completely unaware that this is an IPv6 enabled application. Looking around the market suggests that the application development community in the U.S. will need a wake-up call to avoid losing the Internet application leadership role to Asia.

Domestic issues

With respect to domestic efforts, the Moonv6 project is a good start, though currently it is not broadly available to additional participants. The 6net project ¹⁰ is a European predecessor to Moonv6, and its 9,5M €funding is comparable in scale to the EC Euro6IX project ¹¹. 6net is focused on documenting European deployment and management issues, while Moonv6 to date has been primarily focused as a wide-area interoperability test lab for the DoD. Discussions are ongoing regarding establishing a document archive that would assist federal civilian agencies, state and local governments, academia, and the private sector in their deployment planning. To

⁹ <http://verio.com/about/newsroom/pr/index>

¹⁰ <http://www.6net.org>

¹¹ <http://www.euro6ix.org/main/index.php>

date, the normal leading R&D parts of the federal government are operating or testing IPv6, but this point the rest of the government appears to be waiting for commercial availability. While this path makes sense for the majority of evolutionary services, IPv6 represents enough of a step that concerted cross-agency advanced planning may be required to avoid chaos. In the short term, as long as agencies do not remove their existing IPv4 capability, there is no immediate impact on cross-agency interactions. Available expertise is or will be a major factor as individual agencies decide whether and at what pace to deploy IPv6. Given this, a document repository like the one from 6net will be one of the most valuable results from the Moonv6 project. Rural deployments will generally rely on expertise from nearby universities, as they do for other technologies. Without a widely distributed research funding model, universities will lack the expertise to provide that assistance. An area of research is geographic addressing¹², where at least one suburban community is investigating the approach to see if it could be used as a basis for replacing street addresses when dealing with government agencies.

The current allocation policies of the RIR's restrict direct IPv6 allocations to ISPs, with the expectation that other market segments will acquire allocations from their provider. There is an open question about how local policy and business practice will effect small and mid-sized business allocations. The current RIR policy measures the management efficiency of each ISP in terms of how many /48s they have allocated to customers. Since there is no benefit to the ISP in being more efficient than the /48 metric, it is expected that all IT organizations will be able to get at least that much space (80 bits available to the end site). In practical terms, only the Fortune 50 would find a /48 constraining, but if those organizations ask for more they should be able to get it using a similar basis as the ISPs do to receive additional allocations. If the policy were changed to allow direct allocation to non-ISPs, verifying whatever qualifications are in a new policy might require scaling up the RIR staff.

International issues

The tax incentive program in Japan clearly spurred a faster deployment than might otherwise have happened. Acting as a catalyst, this program has moved their market in a clear direction, where application innovation can proceed unfettered by uncertainty. Recent visits to Japan have found many vendors thinking outside the box, investigating the open potential of new market opportunities. For example, during the IETF held in Yokohama¹³, there was a demonstration application to provide finer granularity to the rainfall patterns than was possible using existing radar technology. By outfitting the taxis with IPv6 enabled sensors on the windshield wiper speed controls, they were able to average in the human perception of rain density based on the data sent back to the weather center. Something this simple shows the untapped potential of the larger address space.

Last fall on a visit to Beijing, every carrier was demanding to move beyond IPv4 to IPv6-only networks in the short term. From these discussions it was very clear they were feeling political pressure to establish China as a technology leader in a new frontier. Then this week in Seoul, several vendors were running demonstrations of VoIP and home monitoring products that work using IPv6, where their IPv4 counter part failed in the market due to the widespread existence of NAT.

In terms of their GNP governments around the world do not appear to be devoting substantial funds toward IPv6, but as compared with the U.S., they clearly lead. Targeted

¹² <http://www.ietf.org/internet-drafts/draft-hain-ipv6-pi-addr-06.txt>

¹³ <http://www.ietf.org/proceedings/02jul/index.html>

expenditures with long term payback appear to be the typical strategy. The global balance between government initiatives has served to shift early momentum and development/design decisions out of the U.S. Japan is establishing the standard for customer address management, and provider edge demarcation services. Through efforts like their appli-contest¹⁴ they are targeting a clear leadership role in the emerging ecommerce economy. Through the 6net and Euro6IX projects, the EC is establishing the baseline in operational practice for the IPv6 Internet. By comparison, other than Moonv6 the U.S. testbeds have unnecessarily emphasized forwarding speed and completely ignored the more global issues of operational management, applications, and customer services.

(3) Economic, technical and other barriers to deployment of IPv6

Deployment strategy

Overall a successful deployment strategy will mirror the evolution of transportation from foot-paths, to wagons, to rail, to auto, to aircraft. Each of these progressions was done in parallel, allowing each mode to coexist, and natural market forces to drive preferences. In particular the set of transition and deployment technologies defined by the IETF are targeted at decoupling deployment dependencies as much as possible. This allows each organization or individual the opportunity to define a local deployment strategy. At the same time, by leaving IPv4 in place until there is no interest in keeping it elevates the issues that might arise from interoperation of the protocols. Preferring IPv6 whenever available at both ends ensures that IPv4 doesn't stay any longer than necessary, while preferring native over transition techniques makes sure the transition completes. The traditional chicken & egg issue with respect to routing services or applications is resolved by use of tunneling. In the context of roads, the analogy would be building a freeway in parallel with, rather than disrupt the current two-lane dirt path right-of-way. Once the new freeway is open, people will individually choose to move over to it as their vehicles become capable of the higher minimum performance level. Shutting down the older path only becomes necessary when its usage/maintenance-cost ratio no longer justifies keeping it around. Encouraging vehicle makers to get higher performance out the door before the new freeway is open requires simultaneously supporting both environments. While this approach does not get the sleek highly optimized capabilities out on day one, it does avoid a lag between availability of service and use of that service.

Specifically in the context of IPv6 deployment, the tunneling technologies of 6to4, isatap, & teredo allow end system stacks to present an IPv6 network upward to applications, despite the lack of routed service. While less than elegant, this approach allows application developers to focus on the application rather than the underlying network. At the same time, service providers will have the opportunity to measure actual usage (by watching special IPv4 packets where the next protocol is IPv6), enabling them to time their expenditures and service introduction to match their local return on investment strategies. Since tunneling does represent a cost in every packet (an additional 20 bytes), service providers will be encouraged to deploy a native IPv6 service at the point where recouping the ~10% overhead will allow them to delay upgrading service components. Competing local access providers might use the availability of native routed service as a differentiator.

The primary factor that will influence an organization's decision is cost. Cost is a driver either to deploy or delay, and as such is a very local question. There are many components to cost, and each has to be individually evaluated. The key component gating the other cost factors is availability of IPv6 enabled applications in support of the critical task. Without IPv6 in the

¹⁴ <http://www.v6pc.jp/apc/en/invitation.html>

current-use application, the other factors of cost are usually irrelevant. This only changes when the other costs become high enough to offset the cost of switching or developing new applications. When new challenges are taken on, the costs are biased by application developer training and availability of sufficient addresses to accomplish the task. Some applications, such as real time monitoring of inventory or livestock, will require substantial address space and clearly offset any developer retraining. Other applications will be more affected by the cost of operations, so the short term development costs need to be balanced against the long term cost of managing a complex IPv4/NAT network environment. In these cases, the short-sighted rarely see the value. When all the hidden costs of maintaining an IPv4 network are exposed (such as regularly readdressing subnets, or correlating addresses across a NAT for diagnostics), people will have an opportunity to balance those against the cost of recertifying their environment based on IPv6. As long as the strategy of acquiring IPv6 capability via normal life-cycle refreshment is followed, the costs of switching will be limited to retraining the management staff and reconfiguring any local customizations or tools. Additional cost factors biasing the decision are things like ability to get additional addresses fast enough to meet growth needs, interactions with suppliers/distributors/peers in other parts of the world where allocation policies may differ, and the operational burden for any NAT work-arounds necessary to deploy a new application.

An area of continuing theoretical concern is the impact of the larger address on latency sensitive services. There are clear engineering trade-offs to be made between link utilization and cpu cycles needed to manage compression. The IPv6 header compresses better than the IPv4 header because the options and other variable values were moved out of the base header. Yes the IPv6 header starts 20 bytes larger, but when it makes sense over low capacity links, the compressed header is actually smaller. Engineering trade-off evaluations about the performance penalty of compressing the IPv6 header must include the IPv4 performance penalty of NAT traversal schemes and state management, as well as the lost opportunity for the collection of applications that are simply not possible when the end points are not uniquely identifiable.

IPv6 does not require new routing technologies. At the same time, there are operational challenges today that might lead to new routing protocols. The difference is that any changes would have nothing to do with the version of IP beyond the ability to scale to a much larger network using IPv6. In the short term it is most likely that the default-free-zone of IPv6 routes exchanged between providers will be significantly smaller than the ~150,000 entries in IPv4, potentially resulting in less than 20,000 entries. This is partially due to simple aggregation as organizations are able to acquire a single large block corresponding to the size of their collection of historically disparate IPv4 prefixes. The rest will be due to the current policy of strict CIDR and provider based allocations. As enterprise organizations push back against that policy the size of the IPv6 routing table will grow.

Interoperation between the protocols is a function many people falsely assume will need to be there from the start. As long as network managers leave IPv4 in place, and use IPv6 whenever both ends are capable, there will be little need for interoperation. The only time interaction becomes truly necessary is when an IPv6-only system needs to contact an IPv4-only system. This will clearly happen once a significant part of the market has moved, so vendors start creating IPv6-only appliances, while the more expensive central resource servers have yet to reach the end of their natural life-cycle. For these cases translation technologies exist, and focusing them on specific servers limits their operational impact.

In summary the overall strategy should be to start from the most general transition technology (dual-stack) and work to the most restrictive (application layer translation), stopping as soon as one fits. There is ongoing work to identify which transition technologies or

combinations apply in particular deployment scenarios. In general 'interoperation' between the protocols is problematic, but in targeted situations it is necessary.

Barriers

Technical

Just about any effort to specifically shift to IPv6, over a short period of time in the near term will be more expensive than later. The transition technologies were specifically designed to enable a prolonged overlap, and to minimize the deployment and operational interdependencies, thereby aligning replacement decisions with normal life-cycle timeframes. As far as limitations being a barrier, there are inherent limitations in any system, the context really depends on what each person is trying to accomplish. The engineering that went into IPv6 was specifically targeted to incorporate lessons-learned from the things people were trying to do with IPv4. At the same time, all changes had to be made with minimal impact to other layers of the architecture. In the context of the current Internet architecture, IPv6 functionally replaces IPv4. In the context of other potential architectures, IPv6 has the inherent limitation that it is only designed to fit within and expand the current Internet. The perceived technical limitations that have been raised to date are the lack of a simple solution to the complex problem when sites connect to multiple providers, and the operational difference between the provided multiple addresses per host method vs. current IPv4 practice. Ongoing work to deal with the problems of multiple connections has focused on alternative architectures for the entire Internet. In any case, IPv6 offers the same technique for the multiple connection issue as IPv4, plus a new approach that may work in some environments.

To the specific question of practicality; there are cases where IPv6 is very practical today. It is running as a regular service in Tokyo, where peer-to-peer gaming & game development is a leading application. Another case where it is practical in the U.S. is for the mobile phone service providers; where they can get IPv4 allocations from ARIN, but the rate of allocation by policy is inadequate to support the subscriber growth rates necessary to make their business plan viable. There will be cases where a shift in the short term would be impractical. In particular any hardware or software that has been built with a 32 bit IPv4 address assumption will need to be replaced to support IPv6 at the same performance level. As noted earlier, the transition technologies defined by the IETF allow this to happen with minimal inter-dependencies.

The point of either version of IP is to isolate the end-to-end packet layer from the morass of transmission technologies that are used below, meaning that no modifications are required. In particular, IPv6 will run just as well over IPv4 as it will any of the technologies listed in this question (or even barbed-wire), and in doing so will treat the legacy IPv4 Internet in the same manner as a global frame-relay service. In addition, the transport (as the IETF uses the term) protocol that rides over IP usually does what it can to take full advantage of any available resource below IP. IPv6 has changed the default expectation of the minimum maximum-transmission-unit (MTU) of any link. Thus the average size of a packet might increase (and with a working PTMU discovery might significantly increase). Since efficiency of utilization on the underlying media is a function of any header processing limitations, as well as the number of bytes associated with each header, the result will be very situation dependent. The obvious short term improvement will in those networks which make heavy use of end-to-end IP options. Since the end-to-end options are moved out of the base IPv6 header, routers along the path will be able to ignore them. This will improve forwarding performance relative to the IPv4 situation where end-to-end & hop-by-hop options are mixed, meaning any option requires special handling by all nodes along the path.

The necessary transport protocol changes were documented in the base IPv6 protocol document, IETF RFC 1883 (December 1995), and its updates. Any part of a system of interacting applications that reaches down and grabs IP addresses will need to be fixed or replaced. In particular applications that have a fixed allocation of 32 bits for address resolution responses will need to be updated to allow for 128 bit addresses. At this point we will not really be able to tell whether and to what extent the transport layers might need to be modified the creative energy of the development community is fully focused on the potential offered by a return to globally consistent addressing. In the short term, the design goals remain focused on minimizing the disruption while introducing a larger address space. Over time, we will see continued evolution, just as there has been with the previous versions of the IP suite over the last 25 years.

As far as efforts to safeguard the integrity and security of communications traffic, or limit government's ability to protect legitimate security and law enforcement interests, impacts would be limited to the degree that enforcement infrastructure would need to understand, and endpoints would need to protect both protocols equally. The fact that one protocol could be carried within the other is not limited to transition, because either protocol could be carried within another header of the same version. Since the enforcement infrastructure needs to be capable of dealing with tunneling in the normal case, the tunneling transition mechanisms don't introduce any new security concerns. Translation mechanisms could be used to obscure the origin or destination of traffic, but this is the security concern introduced by IPv4/NAT, so it is not new.

Economic

In many cases, the deployment situations are so unique that any cost estimates would be limited in value to a specific situation. Economic impact will be minimal as long as normal life-cycle replacements are the path for IPv6 deployment. In general, the cost to provide IPv6 capable replacements will be whatever it cost to put the 32 bit specific hardware in place the first time. This is because the marginal equipment costs for handling the larger addresses are very low. CPE update costs are technology dependent. Those devices which are strictly layer-2 bridges will not need to be changed. Any device which looks at the IP layer will need to understand IPv6 as well as IPv4. Actual cost will depend on how much is software based vs. 'burned' into hardware. In addition to training, and updating any customized tools, updating test suites will be an identifiable cost. Training is likely to be the most significant, as universities are not turning out IPv6 aware network engineers. This means not only existing staff will need retraining, but new graduates will also need specific training until the curriculum catches up.

VoIP is one instance of the applications that would benefit from use of IPv6. 3degrees as 'a shared environment with multiple media streams' is another instance. VoIP will not directly drive IPv6 adoption without the accompanying capability of service guarantees. Unless/until VoIP is better, cheaper, or both than circuit switched voice, deployment will be limited to the adventurous. Since the QoS capabilities of IPv6 start at exactly the same place as IPv4, the only value IPv6 adds is the ability to make the VoIP application work at all between arbitrary end points. For two endpoints that have public IPv4 addresses, but use IP options, IPv6 may or may not improve performance depending on link speeds and cpu capacity of the routers along the path. For the more typical private IPv4 addresses behind a NAT, IPv6 improves performance by simply making it possible for the application to uniquely identify each target device at all.

As with any shift in technology, some manufacturers will continue to produce IP4-only products, until their customers switch to a supplier that has lower product costs or lower operational complexity costs by using IPv6. This will not be a particular problem for the

purchasers of those products until a significant number of IPv6-only services exist. Even then, translation technologies exist for situation dependent use. The only market condition that would persuade manufacturers to cease offering IPv4 equipment is lack of customer demand. There is a vast difference between IPv4 capable & IPv4-only capable equipment. The successful suppliers will offer simultaneous IPv6/IPv4 capability until customers are no longer willing to pay the costs for supporting IPv4. Given the independent nature of deployment decisions, IPv4 support in equipment will be required for years to come. At the same time, software may need to have changes, but again this is situation dependent. For example applications that are completely unaware of the IP layer will work without changes, as long as the stack below them keeps that level of transparency. Most software that is aware of the IP layer will only require the ability to handle the larger packets. The only time a software update is inadequate is when there is 32-bit limited task specific hardware.

Internet services will have to be modified to make them compatible with IPv6 transmission, but services like DNS provide responses that are version independent from the transmission protocol (ie: AAAA response carried over IPv4, or A response carried over IPv6). Software that already includes IPv6 in current versions, like DNS, will not be any more expensive than upgrading to the current version. Out-of-cycle changes can be directly attributed to IPv6 deployment costs, but are only necessary in support of a specific application that will not work any other way. Since the quality of software is a local issue, the magnitude to effort depends on the situation and historical attention to software development. Software that is written with a significant number of implicit assumptions about 32 bit address values will require significant changes at an equally significant cost to find them all. Software with clean structures for explicitly passing around addresses will only require minor changes to those handlers & interfaces.

Fortunately in the overall scheme of things there is very little hardware that is truly IP aware. Given this and the set of transition technologies from the IETF, it is very likely that order of deployment will not matter. There will be a few situations where tunneling is not possible and very high performance hardware is required before an application can be useful. In those situations, hardware deployment will have to occur first. Other than those though, the target was to localize the timing and cost decision as to when IPv6 capability is acquired, and/or enabled. There might be business reasons for updating the edge applications first, then using measured traffic volume to determine when to upgrade hardware. The tunneling technologies allow this scenario, such that the applications can acquire the benefits of address uniqueness in IPv6, without waiting for someone to build a business case to upgrade the interconnecting networks. The tunneling approach can also be used to enable network pockets before the entire path. Converting vs. incorporating are different approaches and result in different deployment problems. From a complete conversion perspective, synchronization is probably the biggest problem. Since integration with simultaneous support is available, the problems there will be limited to any 32 bit specific code or hardware.

Training costs will depend on how well trained the current staff is (for example can they deal with hexadecimal? the number of recent computer science graduates that can't has been somewhat amazing), as well as the size of the staff, and how well disciplines are aligned with the network stack layers. For the day-to-day operations staff that only worries about router configurations, the costs might be limited to how to deal with colon delimited hex rather than dotted decimal addresses. For host administrators, the costs will be that plus dealing with potential differences in how addresses are configured on each system, as well as how many addresses are assigned to any one interface. The list goes on, but until the University system starts turning out graduates that are conversant in IPv6, there will be a direct cost to every network with

a management staff, and every consulting organization will have to seek out training. Training is just one of the overall cost factors in any decision. That one time event may or may not tip the balance given the overall costs of maintaining long term work-arounds to complex & brittle IPv4/NAT networks. Some organizations may look at the training as a career advancement opportunity and bury the explicit cost in an overall staff development program. Others may look at it as a significant identifiable hurdle to be avoided because they never bother to add up the complexity induced costs over multiple years.

Opportunity costs vary completely by perceived goals. For example, parts of the military segment (with a substantial IPv4 address pool) will be looking to continue their current tasks at reduced cost, so IPv4 may present a better opportunity cost. Other parts of the military segment will be looking at emerging scenarios including the number of independently addressable devices necessary to accomplish their task, and then recognizing that IPv6 is their only path. In these cases, opportunity cost with respect to a non-existent IPv4 alternative doesn't make sense. For retail establishments looking to track millions of inventory items on a global scale, the opportunity cost is based on the trade-off between continuing with manual efforts vs. moving to automated always-on technologies (like rf-id). Waiting simply means more delays, errors, and salary for manual counting, vs. the one-time cost of an automated system which can provide current accounting without the substantial staff salary & benefit costs.

The extent that the transition path of the U.S., relative to the rest of the world, will influence costs and prices of IPv6 equipment, services, and applications depends on how active DoC becomes in the process. Currently the U.S. is behind the curve for IPv6 deployments, and stands to lose its lead in overall Internet technology development. As a component part of the current Internet, the U.S. market is significant, but at the same time is small with respect to the untapped potential areas of growth. When the rest of the world moves to IPv6 (because that is their only path to adequate address space) while the U.S. delays, the overall costs to the U.S. for applications will be higher due to the need to operate NAT work-arounds. At the same time, as with any technology costs will decrease per instance as the development costs get spread over more instances. Since it takes time to deploy each technology instance, decreasing costs over time is a natural outcome. Overall costs to the installed base would have been rising over that period to deal with the increasing complexity of operating the work-arounds to keep the existing IPv4/NAT technology going. The visionary will recognize the investment return in switching away from the dead-end technology well before it becomes a hard requirement, despite a potentially lower cost for waiting.

R&D costs for some cases will be absorbed by the international market, but in balance, the R&D costs and security problems related to IPv4/NAT work-arounds will quickly swamp any costs related to moving to IPv6. Directing people to wait for the rest of the world to move simply means that significant development and decisions for the next generation Internet will be done outside the U.S. If that happens, the DoC will have been complicit in degrading the overall U.S. economy by handing the economic engine, with high-value jobs in advance application & service development, to the rest of the world. Late entry into global IPv6 markets by U.S. firms will have a significant long-term negative effect on economic performance. Government actions in Japan & China show they already see the lethargy in the U.S. as an opportunity to seize economic leadership. The applications being developed there will not only be cheaper than the IPv4/NAT based ones in the U.S., they will have more creative freedom as developers are free to focus on what people will pay for, vs. whatever can be hacked together as barely passable. Once the leadership in Internet services and applications is lost, it will be virtually impossible to get it back. There is always opportunity to use lessons-learned to avoid problems. At the same time, the resulting fixes will be based on the needs of those first-adopters, and may not quite fit the needs

of the late-comers. The opportunity cost of waiting to deploy is much more about who controls the next generation Internet technology development, than it is about the actual costs for deployment.

The impact of slow IPv6 deployment on the development of native IPv6 applications is virtually zero. At this time it only makes sense to build IPv6-only applications when the target environment is self-contained and doesn't need to interact with any existing IPv4-only node. At the same time, the cost to the developer for supporting both IPv4 & IPv6 is negligible. The strategy of an extended overlap period will allow individuals to move at their own pace, without impacting existing operations. During this period, any new applications that can run self-contained, will find that it is simply cheaper to leave IPv4 out altogether. The only negative impact of slow development & deployment of IPv6-only applications is on mind-set. People tend to be myopic and want to see 'the world is going IPv6-only' before they bother to worry about making a change. Unfortunately by the time they wake up, the world will have long since left them behind.

The incremental costs resulting from operating IPv6 and IPv4 concurrently are no more than the incremental costs of running IPv4 along side its predecessors were in the past. The dual deployment approach has been used successfully over many protocols to get us to a unified IPv4. Unfortunately we now find people are throwing up their hands over the increased cost of adding IPv6. Yes there is an interim extra cost to simultaneous protocol support, but that will happen at some point anyway. The real question is how much will it cost to train staff before a generational turn over provides graduates already up to speed, vs. how much will the complexity and lack of some applications impact the ability to get the fundamental job done? As soon as the time savings to the non-network staff exceeds the cost of network staff training, the actual incremental cost for running both protocols together will be irrelevant.

Some of the transition tools are clearly intended for end-game scenarios (either early or late), but the basics of dual-stack & 6to4 tunneling will scale well. There might be a bubble if half of the nodes are tunneling, while half are not able to do that, but that would only happen if people turn off the capability of tunneling after they get a native service. Since IPv6 does not limit one to an either/or scenario, it makes more sense to leave both native and tunneling up and minimize the demand on tunnel terminating relay routers. Following the preference rules in IETF RFC 3484 will lead to use of native service when available at both ends.

Almost all equipment is capable of more than IPv4-only already, so one more protocol does not create a significant cost. Handset manufacturers complained early on about the memory impact of both protocol stacks, but quickly realized the ~ 10% overhead was easily offset by the ongoing operational complexity of trying to manage translation servers. The impact of the transition technologies on operational networks depends on the situation. Target functions like 6to4 routers can be placed in dedicated boxes, and otherwise have no effect on any existing services. Edge functions like NAT-PT if placed in the middle of a network can have a dramatic impact by skewing the overall perception about what the network is capable of. There are different technologies to address different deployment scenarios. There are tunneling techniques for both public & private network deployments, as well as one that works in the presence of NAT. There are stateful & stateless translators that work at each layer in the stack to mask what is happening below them. These translators are very niche specific and only need to be considered for the cases where nothing else makes sense in the local environment.

The benefit of dual stack is that it allows IPv4-only apps to continue working (at least as well as they currently do), during the introduction of IPv6 into the environment. By preferring

IPv6¹⁵, updated applications will transition at their own pace of deployment. Once IPv4 is no longer used, it can be removed from the environment. The cost of doing this is maintaining and managing both protocols during the overlap (as noted earlier, this is how we got from X.25, Decnet, Appletalk, SNA, etc. to IPv4). When parts of the routing infrastructure can't be moved in the same timeframe as the edges, or other routers, tunneling makes sense. There are tunneling approaches varying between manual configuration and full automation, with high operational costs on the manual end, and low costs on the automated end. The tunnel-broker approach sits in between with the ISP end automated, while the customer end needs manual establishment. In any case, these approaches present the appearance of an IPv6 network upward, while treating the IPv4 network as a global non-broadcast substrate (much the way IPv4 runs over frame-relay or ATM networks today). The costs for tunnels vary by how much automation there is, and how much impact the extra 20 bytes impacts overall capacity or latency. Configured tunnels present what looks like managed circuits between the tunnel ends, so and for all operational purposes it can be treated as a circuit. Tunnel broker services automate the ISP end and delegate prefixes based on the authenticated user. The tunnel state is usually more dynamic than a configured tunnel, but may be based on static assignment of a prefix to any individual customer. Automated tunnels like 6to4, isatap, & teredo embed the IPv4 address of the tunnel endpoints into the IPv6 address. This is a simple technique, but has restrictions. 6to4 requires access to public IPv4 addresses on each end, so it won't work through a NAT. isatap allows private IPv4 addresses, but is restricted to use within a private network. teredo works across the public network, when there are NATs in the path, but is restricted to use from an endpoint rather than a router. Other than NAT-PT, the translation technologies are so environment specific they aren't worth the time here. NAT-PT is like IPv4 NAT in that it translates addresses, but it also translates protocol versions on the fly. It has the same limitations as IPv4 NAT in that many applications are broken when the header and contents don't match, as well the inability to demultiplex multiple devices behind a single transport layer port number.

The embedded base of IPv4 equipment and applications function as a barrier that could isolate the U.S. only in terms of mind set. As long as people look at the deployment as an either/or situation between protocol versions, they will get stuck on the size of the installed base. As soon as they look at the ability to simultaneously deploy both, and weigh the cost of losing the Internet development lead, they will get past any issues with installed base. The only real costs are in those places where there is dedicated 32 bit specific hardware, or a 'need' to upgrade out of cycle with normal attrition.

(4) The appropriate role for the U.S. government in the deployment of IPv6

The U.S. government should concern itself with two significant issues with respect to IPv6 deployment; the economic engine of Internet technology leadership, and access to public information.

Economic engine – social & economic welfare
Balancing intervention

In protected markets where geography tends to limit exposure to external forces, a hands-off approach makes sense. As the Internet is a global market, it is appropriate for the U.S. government to maintain balance by asserting as much pressure as other governments do. The goal would be less about driving the market, and more about keeping pace to avoid losing intellectual

¹⁵ <http://www.ietf.org/rfc/rfc3484.txt?number=3484>

leadership as the rest of the world shifts directions. One characteristic of a leader is vision, which in this context translates to the ability to predict pressure points which will change the buy/develop cost point resulting in a shift away from buying U.S. products.

From the questions, it is not clear if the DoC is concerned about global or national social welfare. In any case, if the U.S. leaves the rest of the world to develop IPv6, the rest of the world will become the leader in defining Internet technologies. The resulting economic shift will in turn shift social welfare in a cascading crescendo. Is government intervention needed? Probably to the degree of focusing people on the inevitability of an IPv6 deployment, and balance the playing field against actions of other governments.

There is a perception of a chicken & egg problem, but lack of infrastructure is not really a problem because the transition technologies were defined specifically to break that dependency. The thing that is holding off application developers appears to be lack of faith that there will be customers. IPv6 is plumbing, and rightly typical end users will never demand specific plumbing. The application marketing departments are looking and saying since nobody is demanding IPv6, we won't build it in our products. This short-sighted (typical limited Wall-Street next quarter results focus) approach maximizes short term profits, and works as long as nobody is jumping ahead. We already see evidence that Japan & China are trying to do exactly that, so the lack of local capability means that all production shifts out of the U.S. as soon as demand for a cool-new-application hits the streets. The one thing the DoD statement did was put a visible customer out there that was specifically demanding IPv6 plumbing. The DoC could help with the market uncertainty by establishing a clear direction with reasonable timeframes, then providing the information repository to facilitate training.

From a regulatory standpoint, IPv6 is similar to other technologies in that FM & HDTV are about managing the scarce spectrum resource, while IPv6 is about the scarcity of the IPv4 address resource. It is also similar in that there is a massive installed base with little short term market incentive to change direction. The similarity ends there as there are local national standards that retain the characteristic of a local market place for FM & HDTV, while the Internet has no borders, meaning there is one global market. Left alone, other governments will set the direction of that market, and the U.S. will be left to follow.

Access to public information

The other role for the government is to facilitate access to public information. This includes both direct government information, as well as general Internet resources accessed via public access facilities such as libraries. In making government information accessible requires the provision of IPv6 on government information servers. Other than the DoD, the U.S. agencies are not demanding IPv6 capability in the products and services they buy. Since the government represents the single largest consumer of services in the U.S. market, the lack of demand sets a tone which in turn impedes application development, and in turn deployment. While the cost of enabling the servers is a direct funding effort, it also has an indirect side effect of funding training for the contractors that typically conduct the actual hands-on part of the deployment. This expansion of knowledge base will result in another substantial return.

Market intervention Need

In general it would be true that IPv6 applications and ISP routing services would be interdependent, but the transition technologies were specifically developed to break that dependence. Yet if people insist on trying to operate IPv6-only networks that need access to the IPv4-only network, they will incur an interoperability cost. Maintaining the existing IPv4 for access to the legacy systems, removes that interoperability cost. The only impact this has on the chicken & egg problem is the perception that the potential market for IPv6 enabled applications will be limited. Since the transition technologies specifically decouple the infrastructure from the end system deployment timeframes, the potential application market is limited to end system deployments. Since IPv6 is a native part of Windows XP & later, and *nix., it is simply available for the majority of application developers today.

Form

Market encouragement might take the form of short term tax relief (as Japan did for 2 years), or government procurement expectations (as DoD is currently doing). It could also take the form of a National Security mandate (making the job of lawful intercept easier, reducing the attack profile of any given subnet, enabling ad-hoc event scene networks), but these stick approaches should be limited until a carrot approach has been tried. To the specific discussion, 'We therefore seek comment on whether any firm or firms have monopoly power for IPv6 products and services, and how the exercise of such monopoly power will affect IPv6 deployment in the United States.'; There are no software monopolies, as software is just an instantiation of ideas. Ideas are not limited to any firm or group, nor are they limited to any country. As such impediments to deployment have more to do with fears (something governments typically foster to feed off of) and lack of training, than they do with any point of control. While economists will dismiss early entrant power as a long term problem, they will also recognize that when backed by foreign governments, these early entrants are not on a level playing field. Since the Internet is a single global market, the only ways to offset the outside influences are to be the idea leader, or have a comparable background influence. Since idea leadership is being handed offshore due to lack of focus, government intervention to balance the playing field will be required.

As a straight up technology, IPv4 & IPv6 are by design functionally equivalent, so premiums are not really viable. That said, IPv6 builds on IPv4 in that it enables all of what IPv4 does, plus it repairs the damage done by the recent deployment of NAT. In this way it turns back the clock to the early days of the Internet, where the promise of an unimpeded frontier offers vast opportunity for application developers. If deployed in isolation externalities would have an impact. Since IPv6 is designed to be deployed along side IPv4, the magnitude of the decision is offset by the ability to do something that is not possible with IPv4. Since IP is plumbing, and most people never look at the plumbing, the decision is really at the application developer and making something new work, or lowering their development cost. Since there is nothing limiting about IPv6, its value would apply to all applications. It would be hard to argue that client/server applications would benefit from IPv6, unless the server was deployed in the network behind the NAT (family album server). The Internet is a global market; there is one standard that applies everywhere. From that perspective, firms that are too focused on the U.S. market will cede their future to the developers in Asia who are being encouraged to seize the open opportunity for the upper hand. Once idea leadership is lost, the U.S. firms will be playing catch-up & have capabilities dictated to them.

At the end of the day, people are simply trying to accomplish a task as efficiently as possible. The tools they use to accomplish the task are the applications that leverage the power of the global Internet. Historically, Internet application development has been lead by U.S. based companies. The direct intervention we see in Japan, China, and by the EC attempts to wrest that

lead away and move the associated economic engine out of the U.S. Thus the greatest threat and need for intervention is in the application development area. Recently I was approached by a global company that has grown to the limits of the IPv4/NAT approach, yet wants to continue to grow. At this time, as they look around the global Internet landscape, they are forced to look for application products and services outside the U.S. While this is one simple instance, the cascading effect of shifting the economic base will be difficult if not impossible to reverse.

The loss of an economic engine has repercussions in direct tax revenue loss, as well as the indirect impact where ranks of unemployed workers don't have money to acquire goods and services from other parts of the overall economy. No action risks losing the economic engine of Internet application leadership.

There is no need for an information clearinghouse to be an exclusive option, as it should be done in any event. One of the things the government did during the early deployment of IPv4 was to act as an information resource. This underappreciated function of documenting the alternative path is still critical for every new effort. While the open market may be great at finding optimal efficiency on the current path, it is extremely poor at looking beyond tomorrow to avoid box canyons. In this role it is appropriate for the government to step out of the short term concerns and chart out the alternative IPv6 landscape.

At a minimum the government agencies (federal, state, & local) should state their intentions to acquire IPv6 capability in all goods and services by a reasonable date. The DoD put a stake in the ground which had the desired effect on their suppliers, but the short term effect ends up with a pile of exception actions. An appropriate announcement for the rest of the agencies would be to require IPv6 capability in all procurements two years from the date of the announcement, being stringent about allowed exceptions. To the degree possible, the acquisition requirements should focus on the end goal applications, leaving IPv6 to its role as necessary plumbing to accomplish the required task. Any specific action of market intervention would be stating up front that the deployment period will be as soon as possible, rather than waiting until it is as cheap as possible. In this context the short term investment in technology has significant long term returns, both monetary and public-good.

Overall the Fednets & Abilene IPv6 deployments have not provided the community expertise that happened with early deployments of IPv4. Their focus on forwarding performance has not resulted in a sharing of expertise, or development of applications that can leverage that performance. Having recently been at one of the I2 IPv6 WG sessions, I can say the current government support effort is clearly insufficient. While that program is a good introduction to concepts, it is limited to very small groups and doesn't have the resources to go in depth where the attendees could go home and relate the training to their local network.

Matching grants or direct funding of application research requiring the direct peer capabilities of IPv6 is something the government can accomplish. The open market will face explicit training costs until the University system starts turning out graduates trained in IPv6. History shows, that doesn't happen until there is someone doing local research that is available to educate the educators. In addition to expanding the educated base, funded application research will lead to new products and services in the open market.

One specific area of application research would be in efficient K-12 remote learning experiences. While higher-education institutions may have facilities for distance learning, or collaborative research, the majority of K-12 institutions I am aware of are stuck behind NAT, precluding use of these valuable tools. Specific areas of research would be in integrating a 'kid-

friendly' user interface with the complex tools currently used for remote collaboration, and study of curriculum options to effectively present distant information or concepts in the K-12 environment.

In the past I have favored tax incentives targeted to offset similar efforts by other governments. While that may help with infrastructure build-out, focusing on the general commercial infrastructure will not return the biggest bang-for-the-buck. If tax incentives for targeted applications can be constructed, they may still be appropriate. For example a targeted tax incentive for applications and services in support of first responder networks would make sense. In this era where greater interaction between response teams is critical, the priority has to be placed on making the ad-hoc event scene coordination as simple as possible. While IPv4 specialists could build a custom network to support each unique event case, the plug-n-play capabilities designed into IPv6 greatly simplify this effort.

One area where mandates would make sense is intra-government as a supplier or conduit of information. A date should be established as soon as economically reasonable where all government agencies make their public information available via IPv6 for those that have acquired and enabled the new protocol. Going hand-in-hand with that, all government funded libraries should be required to enable access to public information that may be available from IPv6 enabled sites. These mandates should be presented in the context of IPv6 as an enabler to expand access to information via an array of consumer appliances. Once these appliances become widely available, consumers will want to use them to access the public information that the government provides.

Magnitude

The EC has committed 18 M € to the 6net & Euro6IX projects. Japan recently concluded a 2 year tax subsidy through its IPv6 promotion council, but has not released the resulting financial data. China is currently funding construction of an IPv6 research & education infrastructure called CNGI, with projected funding around \$170M.

Given these efforts, it seems reasonable for the U.S. government to provide balance with funding in the range of \$50-100M. The substantial majority of these funds should specifically be targeted toward advanced application research via the University system. Some of the funds might also be targeted at specific First Responder scenario applications.