UNITED STATES DEPARTMENT OF COMMERCE
NATIONAL TELECOMMUNICATIONS AND INFORMATION
ADMINISTRATION (NTIA)
WASHINGTON, D.C. 20230

INTERNET PROTOCOL VERSION 6 (IPv6) WORKSHOP: THE IMPACT OF THE ADOPTION AND DEPLOYMENT OF IPv6 FOR INDUSTRY, THE U.S. GOVERNMENT, AND THE INTERNET ECONOMY

First Amendment Lounge of the National Press Club
519 14th Street, N.W., 13th Floor
Washington, D.C.

Tuesday, September 28th, 2010
9:00 a.m. to 12:30 p.m.

C O N T E N T S

C O N T E N T S

>> MR. STRICKLING: . . . and smart phones to the Internet. We cannot route traffic to and from these devices and without an adequate supply of these addresses, we cannot design cloud computing networks or design the smart grid and as we move to a world where everything can be networked to everything else and all the innovation that will result from that, we have to have plenty of IP addresses.

When the Internet protocol was first developed in the early 1970s, few of the scientists and technologists that were involved at the creation could have predicted what this idea would mean for network deployment and the incredible innovation it would spur.

IPv4, Version 4, was developed in the early 1980s through the work of the technical experts who formed the Internet Engineering Task Force (IETF) and at the time who would have expected that 4.3 billion addresses, as defined by the IPv4 Address Space, would not be enough for future networks.

I asked Vint (Vint Cerf) this morning if he had any idea at the time that it wouldn't be enough and you'll

hear his answer when -- when it's his time to speak on the panel.

But, who could have forecasted the explosive growth of the Internet and the need for billions of devices to be attached to these networks? Well, we're lucky that many of these technical experts did begin to worry about the size of the Address Space at an early enough time and begin to work on the next generation and this new generation of protocol is known as IPv6 and it's good that they did because we now face an exhaustion of IPv4 addresses, but, fortunately, IPv6 will support, I'm told, 340 trillion, trillion, trillion addresses.

So IPv4, four billion, IPv6 340 trillion, trillion, trillion. Hopefully, this will be enough, but we'll see where folks are in 20 years and whether they'll be having the IPv8 or v10 conference, whatever it'll be by that time.

So what does all this actually mean? Well, for consumers not much right now. They can continue to use their existing devices and their IPv4 addresses, even as industry and government upgrade, deploy, and

integrate IPv6, but for industry, particularly the smart phone and router manufacturers, the transport providers, the Internet Service Providers and chief information and technology officers throughout industry, action is needed.

Today, we want to impress upon everyone that IPv6 uptake and adoption is an urgent issue but one that can be successfully handled with good planning and we want to encourage companies to share their best practices on IPv6 adoption so that all businesses will benefit, particularly small- and medium-sized enterprises who perhaps have not been as diligent in pursuing this issue as larger industry has.

So with that in mind today, we're going to have two panels of experts, one from industry and one from the Federal Government, to discuss these issues.

I'm looking forward to a very stimulating and interactive discussion on the importance of deployment of IPv6 and how we in government can work in partnership with industry and other stakeholders to ensure that the technology that underpins the Internet continues to support innovation and growth.

With that, I'd like now to turn the proceedings over to one of our experts that's with us today, Danny McPherson, Vice President for Research and Development at VeriSign, who's going to provide us with an overview of the IP Address Space and set the scene for our discussions today.

Danny is currently leading VeriSign's research in the areas of network security and availability. He has nearly 20 years of experience in the Internet network operations, security, and telecommunications space, having worked at Arbor Networks, Amber Networks, Genuity, Quest, and the U.S. Army Signal Corps.

He's been an active participant in Internet standards since 1996. Currently, he's a member of the Internet Architecture Board (IAB) and the Internet Research Steering Group (IESG) and has worked with the Internet Engineering Task Force, and he also serves on the ICANN Security and Stability Advisory Committee.

So I'm pleased to welcome Danny, this morning's professor, to present this morning's tutorial.

Danny.

[Applause.]

>> MR. McPHERSON: All right. So I get to be the professor with all these folks in the room.

I am truly -- truly honored to be here today, actually, you know, with this distinguished group of panelists. It's kind of -- actually, I was looking at the names there and I had -- I had actually -- I knew everyone. I had worked with everyone in some manner and -- all right.

So -- so how distinguished actually? I've actually worked with five of these folks. As a matter of fact, I worked in some capacity with four or five of the folks on the -- on the panel today. So I am actually honored to be up here today.

You know, with that said, you know, I'm going to give an overview. I know there is a handout that had some -- you know, like a VeriSign logo, but it -- it's an IP primer. I'm going to give an overview of sort of layered protocol architectures and -- and the Internet protocol and sort of what the -- what the crux of the issue is and try and set the stage for these folks to -- to discuss their business and sort of, you know, their -- their needs and -- and issues around this

topic in general.

So with that, if you have the handout and you want to follow along, I think it may help considerably if you're not technical.

So -- so, you know, the Internet architecture, the Internet's a global communications platform. It's sort of a loosely-interconnected network of networks that has no central authority. It's based upon, you know, a 30+ year-old protocol, IP, which was originally developed by the U.S. Government with the name of Robustness to enable multiplex use of existing network fabrics.

IP employs what's referred to as a datagram or a packet-based model. Information exchanged between endpoints is sliced into manageable sizes and wrapped in an IP envelope to comprise an IP packet.

These packets are the entity that's transported across the underlying network, you know, endpoints making only minimal assumptions about what the expectations of the network itself actually is.

The -- you know, the -- this layered model or this layered approach employed by IP allows the adaption or

the evolution of the network infrastructure, so the wires and the satellite and the wireless infrastructure and so forth, to evolve as well as higher-layer applications, like a web browser and e-mail program or some, you know, some application on -- on a smart phone, to -- you know, to evolve without any adaption or any -- any impact on the infrastructure, on the IP network itself.

You can liken this to the Postal Service where a letter dropped in the mail is forwarded based upon the addressing information on the outside of the envelope with no expectation that the intermediate facilities would need to examine the contents of -- of the letter to -- to deliver it or get it to its final destination.

You know, IPv4 has been around for, you know, for over three decades now and has been a key in enabling Internet infrastructure elements and, you know, that's obvious from the success of the Internet we have today.

So, you know, in your slide deck, there's actually -- in the handout that you have, there's a diagram that

talks about the Internet Protocol Model and, you know, it's sort of some of the fundamentals or the principles that -- that have enabled IP to be successful and -- and to enable the Internet that we know today.

You know, these principles sort of start with this packet-based, you know, so these envelopes or these segments of information, this packet-based layered communications protocol, and this layered infrastructure again is desirable to increase the modularity of the protocols and to enable extensibility.

These attributes, coupled with the NTN model wherein systems communicate and adapt independent of the underlying infrastructure, actually, you know, provide the simple and scalable network layer while enabling any connectivity that allows services to evolve and be implemented rapidly on the endpoints themselves.

So with IP, transactions are split into functional layers and IP resides at what's referred to as the Internet work layer, the network layer or just the IP layer. Only IP and higher layers actually operate

end-to-end, from one endpoint to the other endpoint on the network system.

As a result, each device connected to the Internet requires a unique destination or a unique identifier, unique address. You know, there are only, as Assistant Secretary Strickling mentioned, there are only about 4.3 billion of these and when IP was defined over 30 years ago, it seemed like it was a sufficient amount of space. We'll let Vint explain why the shortcoming there.

So anyway, evolution, you know, at the IP layer is complex and expensive because it impacts both network and intermediate devices as well as the end systems and it involves both end-system preparedness and network devices in a more systemic nature.

Oddly, it doesn't actually directly benefit from a lot of the market-driven attributes or effects that, you know, these sort of network externalities that either, you know, selling more network equipment might benefit from or developing applications and finding market-driven users or groups of those applications.

Okay. So sort of the crux of the issue or why are we

here, right? IPv4 address depletion and IPv6. The growth of the Internet has exceeded all expectations. One casualty of this success is that of IPv4 depletion or address space depletion.

IPv4 depletion is not a new problem. It was first discussed over two decades ago.

As a result of this imminent depletion of IPv4 address space, the Internet community responded, developing both short- and long-term solutions. One of the first long-term solutions involved removing the notion of classes in the Internet architecture.

IP's Legacy Classical Model was unnecessarily rigid with fixed-size boundaries sort of negating the sizes of address blocks that could be allocated, enabling little in the way of sort of optimal allocations.

So there were these fixed boundaries and you could only give really big chunks or middle-sized chunks or small chunks. There was no room in the middle and that complicated the optimization of address allocation.

Another technique that was used to optimize address utilization was network address translators or NATs.

NATs have some perceived security benefits and allowed organizations to share a smaller number of Internet-facing addresses with a similar or larger number of internal addresses.

One observation that might be made here is that NATs, because they manipulate the outside header, the envelope information, when the information is in transit, violate the end-to-end principle and, you know, this can become problematic for a number of reasons.

Nonetheless, IPv4 NATs are very widely deployed today in both enterprise and residential environments.

Finally, one more evolution on the shorter-term or the interim side that won't ideally involve long-term is the notion of regional Internet registries (RIRs) that emerged.

These regional Internet registries emerged, you know, to enable resource allocation and help conserve address space through the education of other initiatives as well as by setting policies that encouraged optimal use of a number of resources before new resources could be obtained.

So the regional Internet registries are actually an important piece of the community today where they help with education, optimization, and to enable network providers, network operators to understand how they can best employ IP address or to use IP number resources in our environment.

Finally, you know, in the early '90s, work on the next generation protocol IPv6 began which was finalized in the late '90s. IPv6 varies from IPv4 in several ways, although the most notable is the expansive address space it provides.

IPv6 uses a 128-bit address versus IPv4's 32-bit address space. With each additional bit, you effectively double the size of the address space and, as was mentioned, with IPv6 you get 340 trillion, trillion, trillion addresses as opposed to the four billion that you get through IPv4.

Now, I actually borrowed this analogy from one of our panelists or from the RIRs actually, but, you know, to put that in perspective, if IPv6 addresses were golf balls, there would be enough to encompass the size of the sun in IPv6 addresses. So it's quite a large

address space.

So - on to the next slide, if you are following along in the handout or at all, I guess. So IPv6 and transition co-exist and so, you know, this is really, you know, the crux of the issue, is that IPv4 and IPv6 aren't what's referred to as bits on the wire compatible.

The transition plan from IPv4 to IPv6 was what's known as dual stack. Each system and network device would run both IPv4 and IPv6 and when all the devices were IPv6-enabled, IPv4 would simply be turned off.

This in practice is actually far more difficult to be realized than it was in theory. Any dual stack transition model requires sufficient quantities of both resources, IPv4 addresses and IPv6 addresses, to work and for all the devices to be enabled and reachable on the new stack before the old stack can actually be disabled.

Because of IPv4 depletion, you know, nearing us within a year actually, post-depletion dual stack transition models become more problematic and more complex to implement. In essence, if the transition model

involves everybody running both protocols at the same time and then turning the old protocol off, if everybody can't run the old protocol or all the devices they wish to attach, the old network can't run the protocol, then the dual stack transition model is going to be problematic in that sense and that's sort of one of the primary sticking points with this transition function.

As an example, you know, Y2K, you had both, you know, both the notion of a global systemic, you know, event that had a flag day, if you will, and didn't require this co-existence function which is key. You know, that is, with Y2K, you simply had to locate all the two-digit year data and replace it with four-digit data.

With IP and dual stack, everything must handle IPv6 or IPv4 or both into the foreseeable future which could be quite awhile. IPv4 devices may never be upgraded, for example, because of hardware limitations or, you know, computational capabilities.

On the other hand, IPv6 devices may need to communicate with an IPv4 device. Because there's no

bits on the wire compatibility, network layer protocol translation device, in other words, something that might translate IPv4 to IPv6 and actually make them bits on the wire compatible is really problematic and it also compromises that notion of end-to-end.

So you're -- basically, any time you put a device in the middle of the network that tries to operate and manipulate that envelope header information or even payload contents in that envelope, it can cause brokenness or other problems on the network and from a design perspective that's not a good idea.

So where are we here? So these devices also need to interact with control functions, for example, the domain name system, and an array of other applications that contain IP address objects and must keep state in the network that is discussed, you know, compromises to this sort of end-to-end nature.

So anyway, with that, and, you know, sort of throwing spaghetti at the wall, as I tend to do, which is, you know, why all these folks have me up here, I think, this, you know, IPv4 depletion is near and it's just as predicted.

This isn't a surprise to folks that design the protocols. We're aware of this, you know, 20 years ago and there are a number of short-term stopgaps that were developed and the long-term solution IPv6 which provide sort of the ultimate end goal at this stage. The -- you know, the IPv6 is necessary to remove constraints we've been engineering around for approaching two decades now, constraints being the availability, the wide open availability of IP address space to connect devices to the network.

IPv6 readiness is imperative and preparedness is within budgeting and operational cycles. So if you look at timeframes, which I'm sure a number of the folks on one of the two panels are going to discuss, we're, you know, less than 12 months from the top authority, the Internet assigned numbers allocation authority for handing out Internet number resources, IPv4 addresses, IPv6 addresses, and so forth, will have exhausted or completely depleted their IPv4 address space.

So, you know, finally, I think, you know, one of the things that was interesting when I was preparing for

this and trying to give an architectural overview to people and, you know, provide handouts fortunately that may help to educate them on the topic, I ran across a lot of stuff for folks, you know.

These network externalities, the things that affect people and markets and have this market-driven perspective and say you should invest in this or you should be a leader here or, you know, there's a market opportunity here, you know, sort of a capitalist opportunity here doesn’t exist as obviously at the network layer because these network externalities aren't there.

You know, one thing that I think is interesting and useful about this workshop is that, you know, government being aware of the issues has a number of ways that they can impact or influence change. Two of the least intrusive are certainly articulating the need for change and the need for evolution which is exactly what we're doing here today as well as sort of voting with your dollar, being a leader in the adoption of these technologies and showing why it's important.

So to me, that's, you know, the reason I think it’s important to be here today and increase awareness of this. There is urgency. I don't think that we're headed for a cliff. I think that there are going to be a lot of solutions and, you know, potentially, you know, some problems that are approached as IPv4 depletion approaches, but I think that the Internet community does have sort of these transitional measures in place which again a number of the folks in the room are going to talk about, some of the new short-term solutions that are going to get us across that bridge and provide some of the interoperability, as well.

So with that, I do look forward to, you know, the panel discussions and thank you for the opportunity today.

[Applause.]

>> MR. CHOPRA: Thank you very much, Danny. I think in threes. So I will add a third leg to the stool momentarily, but I want to begin by thanking all of you for being here.

My name's Aneesh Chopra. I'm the President's Chief

Technology Officer and I love my man Larry Strickling for organizing today's event and I'm very grateful for his leadership on this and other very, very important issues on the future of the Internet and its impact on our society economically and for every-day Americans and, frankly, for citizens of the world.

A couple of brief comments and then I'm going to have the rock-stars to the left and right of me provide their perspectives, but I want to set the stage.

In September of last year, the President released a Strategy for American Innovation and in that Strategy he identified the importance of building blocks of innovation, and in that category he referenced the need for advanced IT infrastructure.

Today's topic is just one component of an advanced IT infrastructure for the world. It may seem obscure. We're talking about addresses, but we think it's an enabling infrastructure for a whole range of issues.

We see an ecosystem that's been evolving that is a combination of the supply side, that is those that are providing the equipment, the services that enable the Internet, the demand side, industries that are

understanding the power and potential of this capability and incorporating it into their growth plans, and the rules of the road that are evolving that relate to the supply and the demand coming together and that's where I want to introduce the concept of the third leg before I turn over to my colleagues.

For many of you who are following our policy environment, as you know in the Recovery Act, we've taken some pretty active steps to engage in a voluntary consensus-driven manner with the private sector in areas like the Smart Grid and in areas like healthcare IT, where we've engaged dozens and in some cases hundreds, if not thousands of individuals and organizations to think through what will the Smart Grid in the future look like. What are all the connected devices, be they your refrigerator or anything else, and similarly in healthcare.

If any of you have been to a hospital room, God forbid if you have, you know, for your health and safety, but just think about all the devices today that are essentially not connected today to the Internet, if

you will, and just imagine if every one of those blood pressure cuffs and whatever, you name it, is connected.

>> MR. CERF: It's terrifying.

[Laughter.]

>> MR. CHOPRA: And, so terrifying as it is to Vint, industries are embracing the power and potential of the Internet in ways that many of us are just excited to see the potential of.

So I bring that third leg of the stool which is the role of voluntary consensus standards-based activities and I would highlight as an example of this in January of 2010 when NIST released its protocols or its early round of standards, there was consensus around incorporating IPv6 requirements into the Smart Grid. So I just want to make sure that we didn't lose that third leg of the stool.

So, here's the rules of the road. The rock stars to my right and my left are going to speak for between five and eight minutes to sort of frame for you their perspective on the private sector's challenges here. I'm going to prod them with the following early

questions just to kind of get them to think about these issues as we get into this conversation. As we know, this ecosystem term which is in your hand-out when you framed it, the ecosystem is, as I said earlier, this mixture of supply and demand but you might also think of this in more practical terms, application providers, Internet service providers, and infrastructure stakeholders, be it equipment makers or governance officials and so forth.

What I'd like you to do is self-identify kind of where you are in the ecosystem, however you define it, a little bit of a gut check about, on a scale of 1 to 10, if you will, where you believe your particular component of the ecosystem is relative to preparedness, and so you don't have to indict your individual firm, you can sort of think about the component of the ecosystem you're in, and then obviously the whole purpose of this panel is to think about what we can do to close the gap, 10 being we're at nirvana, 1 being we're uh-oh, and so obviously your thoughts on how we can close it.

This is going to be in your presentations anyway, but

I just wanted to kind of get this framework out there and here's the order.

I'm going to ask John Curran from ARIN to get us kicked off and if you don't mind, John, why don't you rock and roll and we'll hit the ground running?

>> MR. CURRAN: Well, thank you very much. I'm John Curran, President and CEO of ARIN, the American Registry for Internet Numbers. We're the administrator of Internet addresses, what we're talking about today, for this region, Canada, North America, parts of the Caribbean.

I've been involved in the Internet sort of since the beginning, having run two national backbones and a hosting company, and I've been involved in the IETF, including sitting on the task force in the early '90s that realized we were running out of IP addresses and we were going to need a new Internet protocol which we called IPng, which is now beneficially adopted as IPv6. That actually happened in 1999 for people who don't know it. IPv6 has been around awhile.

And in fact, the Regional Internet Registries, the five entities, each one handles part of the global,

administering of numbers. We've actually been assigning IPv6 addresses since 1999.

So the good news is this part of the ecosystem is actually up and running. We've got the protocol people need, we have the administrative infrastructure to issue IPv6 addresses, and it's all good news.

However, I'm one of the instigators of this panel. So it can't all be good news because we wouldn't be having a pretty important panel with a lot of very important people's time unless there was something to think about and what's the important part is the transition.

The IETF, the Internet Engineering Task Force, sort of left the transition part of moving from v4 to v6 as an exercise for the reader. It was not dictated because the IETF's a protocol organization. It doesn't dictate business models and how business models change.

Now, the challenge is that given that there's no one dictating the architecture of the Internet and no one establishing mandatory business models, we actually have a voluntary transition process from IPv4 to IPv6.

I've been informing Internet backbone companies, content providers, enterprises, governments that we're going to be running out of IPv4 addresses. I've been telling people that for 15 years. It's -- it gets kind of mundane. You get to update one number on the slide and give the same slide deck you gave the prior year. You get to do that over and over again. It gets a little boring.

But the reality is we're now down to that final stage (?). We actually anticipate we're now at 94.5 percent of the addresses in IPv4 are fully utilized. We now have a situation where that last five and a half percent we expect to be allocated among the five regions based on demand by next summer. We expect that there will be no addresses available in the RIRs, the Regional Registries, to give to ISPs by the end of 2011.

Now, for perspective, the Internet, for people who haven't realized it, is a remarkable success. It's also growing remarkably fast, even as we speak. When the economies of the globe go up and down, the Internet slows down a little and creeps up but it

never stops growing. It's always growing, and ISPs come to regional registries around the globe every six months to a year and get a block of addresses that let them continue to add customers. Those addresses are then given out to smaller networks, enterprises, governments, not-for-profits, so they can connect to the Internet. It's given out from more websites. These Internet numbers, these IP addresses, are an essential component to let the Internet continue to grow and ISPs, as I said, are coming every six months to a year to get their next block of addresses. We will be unable to meet that requirement in the ARIN region starting some time near the end of 2011. It's that simple.

Now because the transition to IPv6 is not dictated by any particular mechanism, the important part is that ISPs and service providers need to figure out “how do I add new customers when I don't have new IPv4 addresses?” It's not a very hard question to ask. You have the essential input of getting new addresses and yet they won't be available. That's what IPv6 is about.

You should be able to connect up new customers to your hosting service, to your business network, to your ISP, to your government organization, to your enterprise using v6 and still have them get to everything on the Internet.

The challenge? Well, as Danny noted, these aren't compatible protocols. It's very difficult to make them compatible. The IETF really didn't have an option there.

So the net result is that for organizations that start connecting up customers with v6, they're hoping that all the content and all of the websites quickly become reachable by v6. Think of it as another path. Your web servers are reachable by v4. We now need them reachable by IPv4 and IPv6 because the new customers are using IPv6.

It's a simple concept. It's going to be a prolonged transition. We can keep this running. We can keep the Internet running while both of this is in place while we begin to move all the important websites to IPv6 and when we start connecting customers up with v6. We can't do it forever, though. So we have to be

expeditious in how quickly we move all the content over and how quickly the ISPs begin their deployment. Most ISPs are aware of this. They've been informed by ARIN or their regional registry wherever they are in the globe that we're running out and they need a business model based on using v6. We actually -- ARIN sent letters to all of its ISPs and all the organizations with addresses, some 15,000 organizations across the globe, informing them of this several years ago and we've been pedantic in reminding them.

So the fact of the matter is that most of the ISPs know this is happening but it's good to hear from an august panel, some of the service providers here who will be talking about their specific plans.

We also need organizations that serve content to the Internet, content providers to think about how to make that available via IPv6, in addition to IPv4. We'll hear about some of that and the content-networks that help propagate that across the Internet are also an essential component.

So most of the organizations involved in the Internet

ecosystem, from what we can tell, are aware of it. It's now taking those plans and realizing them promptly over the next 18 months that's going to be essential for success.

Thank you.

>> MR. CHOPRA: Thank you. That was very helpful. You didn't touch some of those questions that I threw your way. We'll get them in the Q&A.

Let's turn to Leslie Daigle from the Internet Society to give her perspectives.

>> MS. DAIGLE: Thank you very much, and I'm certainly very happy to be here this morning and I think part of my biggest challenge is going to be to speak only five to eight minutes about this particular topic but I will certainly do my best.

So I'm the Chief Internet Technology Officer at the Internet Society. To get to one of your questions, the Internet Society is a not-for-profit organization that was formed in 1992. It has a number of activities and roles, not the least of which is that it is the organizational home for the Internet Engineering Task Force.

It also does a lot of work in Internet governance public policy work around the globe. We have chapters, we have regional bureaus and chapters around the globe and we also do a lot of education and development.

The Internet Society has a long history in doing training courses, both in the developed world and a lot of work in the developing world, Africa and Latin America and whatnot.

So our perspective and our only agenda in fact is to make sure that the Internet continues to work and to grow as the major platform for innovation. Danny described it earlier as saying that the growth of the Internet has exceeded all expectations.

The issue that faces us here really is a matter for the Internet, the connection of networks. It exceeds the bounds of any given network and what we really need to understand is that the future of the Internet -- to continue that growth pattern exceeding expectations really does need to expand beyond the four billion addresses in the original address space. That's fewer than the number of people on the face of

this planet. That was never going to be enough, not for truly world-supporting technology like the Internet.

Some numbers suggest that in fact there will be 50 billion devices connected to the Internet by 2020 and certainly if they are to be connected to the Internet the way we connect to the Internet today, they're each going to need their own IP address and IPv6 is truly the only answer there.

Network address translators, as Danny mentioned, they are fine for certain purposes. They are certainly going to feature in our collective Internet future, but they are, in essence, a retrospective solution. They can only be made to support the uses that we already have for the Internet. They rely on making accommodations and are based on expectations of what users do and do not want to do.

That kind of environment stifles opportunities for innovations. It makes it harder to develop new applications and services that will be usable by all users of the Internet the globe-round. So the technology here may be obscure to Aneesh's point

earlier, but the impact is very real and it certainly will be deeply felt, is deeply felt by people around the globe.

So that's all very well for the Internet and the whole network, but when it comes down to it, changes are made by changes in individual networks, such as the networks of the companies represented here.

So why do individual companies not want to get into a network address translation future? The reality there is that it's the same reason that you don't want to live in an apartment building with a rental apartment. You can't control the noise of your neighbors and you really can't change the actual architecture of the unit you're living in without causing some rather large impacts elsewhere.

The situation is not really dissimilar for some of the solutions for the co-existence of v4 and v6 going forward. They will work as transition mechanisms but just as we all want to save up and own our own property at some point, they really should only be transition mechanisms.

The issues that they bring forward at a business level

are potential impacts in performance and this is true not just for your network but also for content providers or any kind of an e-commerce site, there's potential impact for your accessibility, for the performance of how your website looks to users, depending on how they are accessing it.

So for U.S. industry, there really are important implications here. There certainly are important implications not just here in this country but because the world is your marketplace. So you want to be able to interact with everything going on elsewhere in the world and you certainly want everywhere else in the world to be able to access you.

So since there are certainly parts of the world that are progressing in IPv6, it's -- it is important to understand what they're up to and keep pace. Japan, for instance, has a comprehensive strategy in place to ensure that their content providers and service providers will support IPv6 by early next year. That's certainly a very aggressive approach but it's indicative of the important directions.

The Australian Government, for instance, has got plans

in place to make sure that their systems are available over IPv6 and their functional internal systems over IPv6 by 2012.

It's interesting to note in terms of “what can governments do?”, the Australian Government actually stepped up their plans when the OECD issued a statement from -- a ministerial statement three years ago -- two years ago saying that IPv6 was important. This brought attention to the Australian Government and they stepped up their own plans.

Developing countries need to be brought to the table. These emerging markets are going to depend on IPv6. If you want to work with the emerging markets, then your focus really does need to be on v6 and not just on IPv4.

Some of the current challenges to deploying IPv6, we're still stuck with companies not perceiving a business model, although increasingly there is recognition that this is a question of strategic expense, but there still are perceptions of business risk.

If your perception is that your content is best

provided over the existing IPv4 network, it's very challenging to contemplate making a leap to the growing IPv6 network. So for that reason, I think it's important to look at what is goodness over the course of the next year. What do we need to see happen?

Frankly, at this point where we are at with IPv4 run-out, I think it's important to contemplate an order of magnitude increase in IPv6 traffic over the next 18 months and it really doesn't matter how you measure it, the point of discussion is to focus that it's not just, yeah, we should get around to doing IPv6, we need to stand up and seriously move to IPv6 and demonstrate that we're doing it.

If we have 20 percent IPv6 traffic by December 2011, that is a clear statement that IPv6 networking is real, content providers are providing content, access providers need to provide access, and business concerns can be allayed.

I don't want to get into talking about the feasibility of achieving that metric. It's merely a number to throw out there to have a target to say what are we

actually talking about.

So at this point I think the only -- recognizing that the only way forward is through IPv6, important steps that companies can make are not only to plan for and deploy IPv6 themselves but also be open and make clear statements that that is exactly what they're doing because if it's done -- one of the things that we've observed in our discussions is that everybody's kind of looking around to see what their friends are doing and if everyone is deploying IPv6 in secret essentially, we're not getting the buddy support system in place to understand that, you know, this is perceived as business-critical and serious effort is being made to deploy.

So that would certainly be my ask because I think it's important to understand that staying with the current Internet is not an option. The Internet is changing. As we run out of IPv4 addresses, we will either have an increasingly network address-translated world or we will move to IPv6 and the latter is the path to growth.

To just add a couple points to Danny's comments on why

this is not quite like Y2K, apart from the fact that there's no fixed date or flag day, it's also important to understand that this is not about hunting down obscure little pieces of code that are buried in elevator controllers and other bits of, you know, sensor software to find the little bits of old COBOL remnants, that this actually affects all of us up-front everywhere.

And another important distinction is that many of the systems impacted in Y2K are independent. It is an elevator controller. It's one company's security system. It's a banking system. But by definition, the Internet is really about interconnecting and so failure to move forward on any front will impact all of us because we're all in this together.

Thanks.

>> MR. CHOPRA: Thank you very much, Leslie. What I'd like to do is now turn it over to Jason Livingood to give us perspective of one of those active participants in this transition, Comcast.

So please join us in welcoming Jason.

>> MR. LIVINGOOD: Great. Thank you very much. I

work for Comcast. You asked where we play in the ecosystem. We're a large Internet service provider.

>> MR. CHOPRA: You're answering my questions. You're a good man.

>> MR. LIVINGOOD: And as such, I think we're currently the largest residential ISP in the United States and play a key role in the Internet ecosystem, both in the United States as well as globally.

I would say in terms of the level of preparation of that sector generally, I would say that it's moderate and it could be better.

In terms of IPv6 and what our activities are, however, which is what I'm here to talk about today, we announced production network trials in January of this year and one of the reasons that we did that, going back to the buddy system point, is to raise awareness in the community, to act as a catalyst with other ISPs, enterprise networks, equipment providers, software developers, and so on to say that IPv6 is real, it's coming soon, and that people should begin their own trials, and we hope that we've been successful in that regard and we've seen a number of

other trials announced.

But certainly there are real business drivers here behind this effort. It's very clear that when you add new Internet service provider customers, they demand IP addresses and existing ISP customers always want more of them and so if there's a point in time when Internet Protocol Version 4 addresses run out, that is clearly a problem if you want to continue to add ISP customers.

So IPv6 addresses are important to the future of that business and are going to be very necessary in the relatively near future and so a part of what our trials were also doing was to sort of explore all of the possible transition mechanisms that we could use, how would we get our customer base to IPv6 and so on, and a key part of that is certainly mitigating any risk related to these technologies and to get out there as soon as possible, find any problems early, if there are problems, to fix them while we still had time.

And it was also to start to prepare our access network which was the last part of our network that wasn't

IPv6-ready for the transition and it's been the case, we've certainly been working for many years on IPv6 and our backbone and all of the core parts of our back-end infrastructure have already been upgraded, and in fact any of our new interconnect partners and pairing partners are adding v6 at the same time that they add v4. So that's been occurring for some time. These trials are more about the end-user and the access network to those homes and so we have a number of things that we're trialing there and I'll mention that in a moment, but we're also doing a lot of and developing a lot of training because if you imagine today, as an ISP customer, you have some question and you need to call customer service and say I can't reach this website and they walk you through paying [for] an IP address, there are whole legions of customer service representatives and technicians and engineers that need to be trained about how to explain to customers and do troubleshooting with IPv6 and we've had great response.

Over 7,000 customers have volunteered to participate and there have been three main technologies that we've

been trialing. One, where you still have the four [IPv4] addresses, and you're trying to get to v6 and have to tunnel or go through a network address translator. The reverse of that, where you run out of v4 addresses and want to access v4, you have to tunnel over v6 to get that, and then, of course, the native dual-stack which others have spoken about here which is really the best scenario in my mind in the near term where you have both v4 and v6 addresses.

And while we've proven so far that all of those technologies are workable and that they will scale in a production network, we have noticed that, of course, you know, as our customers will tell us, speed matters and you will have better access speeds and connectivity to the content you care about if you have direct native dual-stack access over IPv6 and that is a very critical thing to take away if you're a network operator, that that's something that customers will perceive and they will find that important.

In terms of next steps, I think that it's important for content providers to move their content to IPv6. Certainly we perceive that there is a bit of chicken

or egg problem, that why would the content provider move there if the access networks will not, and I think part of our trial was to say that the access networks are moving there and do have plans to do so and so now is the time for content providers to step up and do their part and bring, you know, rich content to the IPv6 world.

But beyond that, you know, I think we need to think about, you know, the Internet beyond where we are today and I think that it's very likely IPv6 addresses will become a major enabler of new applications.

I think a number of people in other fora have talked about the ambient Internet or the Internet of things and the explosion of devices and other things in the home. The ambient Internet, yep, and you mentioned -- exactly. You mentioned the Smart Grid and other things that are driving new applications.

All of those things demand a lot of IP addresses and function better when they have direct Internet connectivity and don't need to go through network address translation devices and certainly mobility and mobile devices, the explosion of those devices will

also drive demand and so it's very likely that the availability of so many more addresses will potentially enable and drive the creation of new applications that none of us here at this table can really conceive of right now and so we think that that's very important to keep in mind and will certainly be a driver.

So I think, you know, in short, in summary here, I think, you know, ISPs are starting to do trials now. We are very serious about the transition to IPv6 and very interested to ensure that, you know, it is seamless for customers. I think, as I talk to the engineers in my staff, I define success for them as no one notices what you do and that would be great, and I think, you know, the support of IPv6 helps to support an open, vibrant, dynamic and growing Internet and that that's incredibly important both for our country as well as for the Internet as a whole around the world.

>> MR. CHOPRA: All right. I'm calling an audible. One question before I move on.

You mentioned the fact that the native dual-stack felt

like it had the better performance. Do you have a way to justify -- just a quick back of the envelope math on cost to deploy these various engineering options? Is it 1:1? Is one twice as expensive as the other? Are they close? Just a little bit of flavor, if you wouldn't mind, on the engineering cost by the bucket, the three buckets you outlined.

>> MR. LIVINGOOD: I think it's close in terms of the initial cost to deploy. I actually think in the long term it's less expensive to have native access because you have fewer middle boxes in the network. It's easier to do troubleshooting with customers and, you know, less -- you know, sort of less encumbrances in the network.

So I think in the long term, it's less expensive to have direct native IPv6 access as opposed to going through a middle box. So I think they're close in the short term and it's cheaper in the long term.

>> MR. CHOPRA: Thank you. All right. Let's turn over to George Conrades from Akamai, [a] very active community.

>> MR. CONRADES: Yeah. Thank you very much, Aneesh,

and I want to thank Jane and John for helping to organize this and the setup by Larry and Danny who got a lot of brush out of the way.

I think this is one of those situations that's both a problem, as has been well discussed so far, and it's a terrific opportunity, given the growth and the innovation potential of the Internet.

I'm reminded by some of these comments about, well, the blood pressure cuffs and so forth, that John -- the late John Sidgmore, back when we were battling it out running ISPs, he headed UniNet, he used to talk about the future of crickets, all of these things chirping, chirping, chirping, and here we are. Actually, he was right. This is -- we're headed also, not just Smart Phones, iPads, TVs, but also these devices that will be using the Internet and communicating with each other and with websites and with consumers. So it's just in greater volume.

To try to knock off Aneesh's setup in terms of where we fit here, Akamai, we think of our role as to make the Internet feasible for robust electronic commerce, rich media, video exchange and business-to-business

applications.

The criteria we use for that are superior performance over the native Internet and, of course, against anybody else trying to do the same thing and greater reliability and so our customers are the content providers. So let me just kind of channel my remarks down that direction this morning.

We have about 3,400 customers who represent on any day about 15-20-30 percent of the Internet's Web traffic. We log over 500 billion IP requests a day on our network. So that puts some dimension on the popularity of these particular websites and, of course, it's global in terms of access to these websites and, not all of them, but most of them are still in the United States.

So obviously one of the challenges we have is working with our content provider enterprises to help them with this transition of IPv4/IPv6 and so our goal, we're -- and this will get to where we're at in this process. Our goal is to have our content providers not have to worry about IPv4/IPv6 transition.

We will take care of that for them so that they don't

necessarily have to make changes to their origin websites to connect to v6 users and we're in -- I like to say we're in beta, others say we're in alpha by the end of this year, and by the middle of next year we should have this capability ready for general availability. So that's where we are and we certainly work with our partners, the carriers, and actually less so with the consumers. We're highly focused on the content providers, the enterprises, and our relationships with our network partners.

Larry mentioned urgency, planning, and sharing, and I thought I'd make a few remarks about that. I actually think this is an issue that should rise to the board-room level among enterprises as well as it is among carriers and, you know, in Y2K, with what we've talked about, we had a date certain, but certainly the issue was risk management and that was an issue that was important to the board of directors and there were quarterly reports.

Anybody that served on a board or served up information to a board knew what you -- every quarter, red-yellow-green, “where are we and are we going to

make it?" Behind that came SARBOX and 404 mitigation for -- you know, to make sure that your financial system were operating correctly and could be verified, audited for their accuracy and the issue there was also risk management because if, in the process of going through your processes, you found that you had a material defect, you had to report that and you had to tell your investors that you had -- you found a material defect in your -- almost always accounting processes and so that was a hair on fire deal in the board room, as well. Every quarter where are we? Red-yellow-green.

The green movement, you know, doesn't have a date certain but it's beginning to find its way into the board-room. There are more and more customers -- enterprises talking about sustainability now. In fact, I'm on the board of a company that is evaluating the possibility of having a sustainability committee. Now boards don't like to have lots and lots of committees, you know, but at that level in this one company where maybe we should have periodic reporting on our efforts to not only how we operate our business

in a sustainable way, but what about our products and the motivations for that are, frankly, investor interests in the subject. Are you really with the green efforts, so to speak, and the other is -- and perhaps the most important, if you didn't like that, maybe you think that's political, but most important is savings. There are cost and expense savings to be gained by really managing your energy usage, etcetera. Well, I think we don't have a date certain on v6, but I think this merits the same kind of attention to risk management. If you -- if the board actually understood what we were talking about today, they might not want to have a report every quarter but I'm pretty sure they'd have a report at least once or twice a year on where do we stand on v6. Do we have enough -- you could just hear them. Do we have enough addresses? Are we able to reach new customers, the points that were recently being made.

So I think Akamai, we certainly talk to the marketing people and we certainly talk to the CIOs.

It's not clear the extent to which we talk to the C-level folks about making this a risk management issue,

but I think it is and I think it would help with the idea of advancing the understanding of this particular problem.

I also agree the government can be a great bully pulpit and the government can be an adopter of v6 and that will help, as well, but unlike some of these other countries where you can kind of mandate -- we were talking about this earlier.

You can kind of mandate the adoption of the v6, this is good old America, you know. This is going to happen for one reason or another and -- but it will, but I think it can be encouraged by having people and enterprises, the content providers focus on what needs to be done to be sure that we are v6-compliant and able to continue to grow.

Now just one thing. I think Leslie made the point about this is not finding COBOL embedded in some program that could affect a date or therefore affect the way a program operates.

Actually, inside of a lot of websites there's some hardwired v4 stuff that tie to a particular server, tie to a particular application, could have problems

with your analytics software and how you analyze the data that you're -- you know, v4 addresses are different obviously than v6 data.

So there is a laundry list of things that a company should take a look at and be reported on if you're going to have this audit by the board as are we v6-ready, let alone do we -- have we committed to it, but are we v6-ready in the way we operate our web infrastructure.

So that's my two cents.

>> MR. CHOPRA: Let me call an audible again on this same point. Do you have such a template drafted inside Akamai; that is, the equivalent of your board report on whether or not Akamai is ready?

>> MR. CONRADES: Well, here's the deal. I'm the chairman and we're going to have one.

[Laughter.]

>> MR. CHOPRA: And --

>> MR. CONRADES: I got so excited about it, I thought, you know, this is really a good idea.

>> MR. CHOPRA: It is a great idea. Audible Part 2.

>> MR. CONRADES: Yeah.

>> MR. CHOPRA: Would such a template be something you would be willing to share and/or collaborate in the development of?

>> MR. CONRADES: Yeah. Absolutely. It's in our best interests, everybody's best interests. I would love to do it.

>> MR. CHOPRA: Audible 3.

>> MR. CONRADES: Yeah.

>> MR. CHOPRA: Just one more.

>> MR. CONRADES: Sure.

>> MR. CHOPRA: Show of hands by people on the panel. Who would participate in thinking through what the -- you signed up, you signed up. All right. We have our first deliverable of the day. That was pretty powerful, man. Thank you.

>> MR. CONRADES: Thank you very much.

>> MR. CHOPRA: Let's -- no pressure, Ram. You gotta come up with something pretty cool.

Ram Mohan, Executive Vice President of Afilias, to give us his perspective.

>> MR. MOHAN: Thank you. Good morning. Afilias is a global leader in advanced registry services, the power

of successful domain registries. We support about 16 million domain registrations, including ".org", the largest DNSSEC assigned registry.

In addition, we also operate a global DNS resolution network and all of our registries currently support IPv4 and native IPv6 and have done so for many years. But expecting an IPv6 server that matches IPv4 is asking a lot. We're quite some distance away. Let me share some facts. The ".org" registry has over 8.6 million domain names registered in its system today. Of these, about 17,000 names have both an IPv4 and IPv6 address in their records.

>> MR. CHOPRA: Wait. 17,000 what?

>> MR. MOHAN: Out of 8.6 million.

>> MR. CHOPRA: Thank you.

>> MR. MOHAN: Okay. And 99 addresses have IP -- 99 names have IPv6 addresses only.

The ".info" registry which we also operate has over seven million domain names. Of these seven million, 58 names have both IPv4 and IPv6 in their records and 25 have IPv6-only.

The ".mobi" registry, which is often used in the

rapidly-growing mobile Internet area, has over 910,000 domain names registered in its system. Of these, no names have both IPv4 and IPv6 records in their records and only two names have IPv6-only addresses.

We also manage the technical systems of 11 national sovereign country code registries with close to one million domain names registered. Of those names, 14 domain names have both IPv4 and IPv6 addresses in their records and six domains have IPv6-only addresses.

In contrast, if you look at ".org", 2.3 million domain names have IPv4-only addresses in them. ".Info" has 300,000 IPv4-only addresses, so on.

So clearly there is some ways to go between the reality of IPv4 depletion and the other reality of IPv4 usage.

Now over the years we've spent time and money migrating all our services to be fully IPv6-compatible. We're not there yet. Let me share some experiences as to what's causing us to get to that point of not being there yet.

Hardware vendors for critical pieces of load-balancing

equipment claim to have support for IPv6. Hold on a second. We procured equipment from various competing vendors and commenced testing. Quickly it became evident that there was a remarkable difference in the way IPv6 flows are processed by the different security appliances currently in the market.

IPv4 packets are processed by dedicated hardware built in the appliance and this allows the appliance to handle filter lists of several thousand entries with little or no impact to performance.

In contrast, IPv6 flows are at the moment processed in software and they do impact the CPU directly and, consequently, it affects the forwarding rate of the appliance as the filter list grows in size.

As a result, we've provisioned a separate set of front-end firewalls dedicated to handle IPv6-only traffic in order to minimize risk on our IPv4 infrastructure. We've also beefed-up our evaluation criteria because support for IPv6 does not equate to equivalent performance.

One of the take-aways for all of the industry is, first and foremost, train your procurement department

because teaching them that best-fit calculations based on IPv4 defaults do not translate in an IPv6 environment.

Let me share a second experience. About the same time that we were looking at load-balancers, we also were evaluating vendors who could offer equipment for rate-limiting traffic inflows into our global networks that would match our current IPv4 policies. As of today, there are no vendors in the market that offer a solution with sufficient level of granularity to match that of our current IPv4 policy.

So for us, the likelihood of having a single appliance with a unique and enforceable policy of rate-limiting for both IPv4 and IPv6 is very unlikely. Obviously that means increased costs, increased management and usability burdens.

Now, one of the questions that -- an issue you had mentioned earlier was, you know, what are some of the things that are doable in [an] IPv6-only environment, innovations shall we say, that, you know, aren't there in the IPv4 environment?

We operate in an environment that is prone to a lot of

probing, gaming, and attacks. Inside the domain name industry, for example, gaining access to the Whois database of a registry with valuable identifying personal information of website address owners is a very profitable, if unlawful in some parts, it's a very profitable act.

We notice that in the industry most Whois rate-limiting systems today limit access based on the number of connections per minute per IPv4 address. Implementation of a similar feature in IPv6 will require refactoring of Whois code everywhere, all the rate-limiting code, to keep track of IPv6 connections across the board.

In addition to this, since the allocation policies for IPv6 customers and providers is done in massive chunks of spaces, it would be relatively easy for a rogue Whois client to implement a system where IPv6 addresses can be rotated in such a rapid amount of time that on the provisioning side, you know, you're just behind the eight ball.

It's innovation but it's something that is pretty unique that could happen in the IPv6 area. I think

that's going to -- this kind of stuff is going to require some policy and business decisions certainly to be done.

So some core principles and some thoughts. First, I think application should be network agnostic. Second, users should never have to know or choose between networks. Hardware, especially network hardware, needs to be IPv6-compatible. You look at routers, firewalls, broadband modems, you know, important components like that.

The good news is that all major operating-system vendors now officially support IPv6 in their operating system releases, but website owners who are content providers need to be accessible via both v4 and v6 and that requires downstream network providers to be similarly so.

I think one of the focuses needs to be on solutions to ensure co-existence between v4 and v6. As I said earlier, v4 isn't going to disappear very quickly and I think dual-stack is perhaps more applicable to service providers than necessarily end users.

If there were IPv4 addresses available for everyone to

run dual-stack, then everyone would have an IPv4 address and there wouldn't be a shortage. So application gateways to me, it seems, are going to be essential but application gateways can cause performance hemorrhages.

So it's not a Y2K problem but I think we have a transition problem and there are some who say tunneling can get us, but I think tunneling can get us only so far. Over time, as Jason is saying, really clarity lies in moving to mainstream IPv6.

Gateways are not good for the same reason tunnels are not good, except I think gateways are worse as a problem.

I mean, tunnels have a downside that you lose some of the benefits of the connectionless network layer but overall my point is that transition technologies could all make this network that we depend upon much worse before it gets better as the end user experience as a requirement for strong performance and responses.

So I see ultimate adoption coming as follows. One, organizations would want to have IPv6 transit so that they can reach ever-growing sets of customers who may

be IPv6-only.

Second, to support this, organizations will need at least their edge equipment to not only be IPv6-compatible but also to have IPv6 routes and addressing configured. Then, they'll want to have their Internet-facing gear to run native IPv6. Once this is done, there could be a significant lag in adoption as companies map out their internal addresses and then map them to external IPv6 addresses. I mean, it's kind of similar to how private addressing works right now.

So, so far to me, it feels like v4 depletion has resulted in some scarcity economics rather than promoting or prompting a massive v6 adoption and a migration. It is my belief that, given the cost of migrating to v6, many organizations are going to add this capacity in kind of a structured manner and they will do so once they can see a direct and tangible benefit. They can obtain rated to their line of business.

I think at the end of the day, economic self-interest is a critical motivating factor, in addition to the

board risk that we're talking about.

>> MR. CHOPRA: So we had a funny back to back. We went from board level dialogue to deep into the engineering stack which is helpful. Audible for you. Question to my friends that -- we're going to get to Verizon in a minute but how -- how well is this sort of learning factored into your trials and your plans for what the engineering tasks are that you're deploying?

In other words, are we rapid prototyping on these concepts in a way that is productive and effective or is this sort of like he's talking in the wind and you guys have already built your plans and you're going here and we've missed the chance to iterate and rapidly prototype? How is that learning factored into the work, if I could just take a second to ask?

>> PANELIST – MR.MOHAN: I think that is definitely factored in. There is definitely a very open feed-back loop in a lot of communication across different players in the community and pointing out the issue of, you know, rate-limiting, that's certainly one of those issues that we've touched on and I think the

only way that we can make folks aware of that is to talk about it in the right communities, you know, where folks are concerned about those things.

>> MR. CHOPRA: Got it. So the question I have for you, just to give the objective perspective of where the industry is, there's a threshold question is/are people engaged? Then the question is that they're engaged and doing something, are they pursuing a path that's least advantageous to the benefit of the ecosystem versus that which is innovation-enabling? I just want to get your gut reaction on the scale of when engaged, the ecosystem's kind of shifting to the less-advantageous path versus the more. Give me a sense of where we are on that spectrum.

>> PANELIST – MR. MOHAN: It seems to me that most of industry's sleepwalking into IPv6 migration.

>> MR. CHOPRA: So when they are engaged, they're engaging in the areas that are least innovation-enabling?

>> PANELIST – MR. MOHAN: That's right.

>> MR. CHOPRA: Okay. All right. Thank you for that sobering comment.

Dr. Nabil Bitar from Verizon's going to give us his perspective. So please, Dr. Bitar.

>> DR. BITAR: Thank you. My name is Nabil Bitar. I work at Verizon. One of my responsibilities is IP Architecture, among other things.

Verizon, as you know, is a global network service provider, providing network solutions as well as service solutions in certain areas. We offer services both in the wireless as well as the fixed wire-line space to both the business/enterprise and the consumer markets.

So by definition, if we list wireless/wire-line business enterprise and consumer, we are in every case deploying an ecosystem involving multiple network provider equipment, multiple application platforms, and we have to make them interwork together and therefore lies the challenge that I'll be talking about later.

What are we doing at Verizon in terms of IPv6? We have to realize that really IPv6 is strategically important to us. As a company, we started early on in IPv6, deploying IPv6 in one of our networks that some

people would known as VPNS+, which is a very high-speed backbone network, serving the government sector in the U.S., and then we evolved from there actually to offer commercial services on our public network and these services to business users and that service started actually about 2007 where we provide dedicated Internet access to our business customers as well as we have IPv6 points and that's evolving, also, into transit, appearing in certain locations.

Along with that, there is an effort where we are -- we have a lab that does certify independent objective vendor equipment for IPv6 as well as security compliance and support which speaks really to the fact that there is inherently a heritage as well as expertise in the IPv6, if you will, evolution in the network and there is a commitment, also, going forward, to really help our enterprise and business customers, whether it's in the government space or the private sector, to chart, strategize, plan and execute the migration strategies to IPv6.

So, really, there is very strong commitment that has shown deployment in the global as well as domestic

networks for IPv6 on the business side as well as to help customers go into that direction.

Now, on the -- we have, as I said, other networks and other services. So, on the wireless space, where it speaks really to what the other speakers have talked about earlier about the IPv4 address exhaust and the problems that it poses.

In the wireless space as opposed to wire-line where we have grown into providing Internet services to our customers, we grew as a market grew. In wireless, we didn't have that opportunity. I mean, it's like really kind of an impulsive function there where we went from almost none to a big explosion where you have to connect Smart Phones and so forth to IP and these devices have to be always connected to IP.

That means you have to be permanently assigned an IP address, always on, so that they could be accessed and could access the Internet and that drove Verizon as well as others you probably have heard in the industry to go to IPv6-way because it's impossible to acquire IP addresses that will address every one of these devices which is larger than the number of phones,

larger than the number of even people because a lot of people have carried and continue to carry and will carry one or two or three devices on them that every one has to be addressed.

Exactly. So there was no option but to go to IPv6 and we've made at least public statements on that, that our devices will be the addressable IPv4 and IPv6. Now the question you asked, why do they have to be dual-stack addressable, the transition to IPv6 is not going to happen overnight. The rest of the Internet will continue. It does today. We will continue to have IPv4 for a long period of time. There are applications that unfortunately have to be built and hardwired to IPv4 that you have to continue to cater to until they are fixed. Until that transition completely happens to IPv4, you still have to provide connectivity to IPv4 and that gives us another challenge that we'll talk about as we talk about the challenges.

Now another sector or another segment we do serve is, as you know, we are also one of the largest ISPs for consumers, providing DSL and what's known in the

industry as fiber services for broadband access. There, we have actually -- I mean as an indication of our planning that's taken place, we have conducted a trial earlier part of this year, back in April, March-April time frame, for FIOS access, the dual stack IPv4/IPv6, that fortunately has concluded successfully and it proved really kind of an architecture that we are targeting for the dual stack for FIOS access. We will continue to plan or chart our plans to really help to bring that to commercial availability, as well. We don't have right now fixed timelines but that's what we're working on. We're planning towers. So we have strong commitment to really move in that direction that we showed in execution as well as in our planning that we're looking to carry forward. Now, the natural question that comes about that we do get asked internally, as probably many people in the audience are asking, why IPv6? We've heard this story before. IPv4 has been around now, IP addresses. It hasn't happened. We've worked around it. We'll continue to work around it. I think this is the cry wolf period, if you will.

Reality is here today. I think we've heard from many and you could monitor, exhaust IP addresses on a daily basis, if you like. It does [is] decrementing pretty fast.

I think the last I looked probably this month, I believe, the availability, as was said earlier, there are about 5 /8's and every /8 gives 24 million addresses available to be assigned to the local registries and that's decrementing. Right now the projection is that will exhaust in mid 2011 and we heard earlier from the ARIN side that as the local registry for the Americas and the Caribbean their capability to assign IPv4 addresses to their askers, if you will, ISPs and large corporations, whatever it is, is going to be inhibited by the end of 2011, I believe. That's what I heard.

So, really, I mean, when somebody asks why IPv6, it's a question that you want to continue to offer services, to add customers, or you want to stop doing that? So, really, the killer is there. I mean, it's really defined. It's not like we're looking for service that you have to go out and define to offer to

users. Really, it's service continuity and service continuity is not just to connect somebody to the Internet or to IP networks but to connect them effectively and efficiently and that lies really with cost as well as the efficient communication.

The IP population, I'd like to call it that way really, is continuing to grow. We talk about wireless phones. We talk about homes being connected to IP. That's growing and that's not going to get any better. Probably it's good for the consumer that it's not getting any better because we're getting more capabilities.

Machine-to-machine is here or about to be here and these devices, everyone is an IP citizen in the IP networks and the IP Internet and they need to talk to each other and not to talk to other devices also connected via the Internet.

So that really -- you could pass them through NAT devices, gateways that [are] in the middle, but that's getting in the way of effective communication, efficient communication, and when I mean efficiency, I

mean also cost efficiency, not just performance. Performance is one dimension but to people who deploy, who pay for infrastructure, it's also about cost because every device you put in the way will incur costs.

The other thing that we think that may happen, that I think if we really go back and look at how the broadband access market had evolved, we've had that for a long time and that came about because of lack of IPv4 addresses that ISP has acquired or the perception that it is.

There are other derivatives that came about from NAT which is considered to be the hiding of devices within the home and kind of the perceived security that NAT does provide, but NAT was not and NAT stands for, for people who don't know that, is really Network Address Translation to enable to translate from one IPv4 address to an IPv4 address in the NAT space.

That really came about from the lack of IPv4 addressing, not really provide security in that way, but it came -- the security came as a product of that.

So now one of the challenges in certain cases when you

say now with IPv6, we have plenty of IPv6 addresses, we could globally and publicly assign every device within the home with IPv6 address, one of the things that come about, oh, but then we're losing security because of that perceived notion that people lived with it.

If you really want to provide security, our view is really you could provide it by having secure firewall which really provide you security. So it's really IPv6 addressing or global assigned devices provide that end-to-end unbroken connectivity that I heard Danny talking about where you don't get in the way of the path of a packet from one endpoint to the other and create hurdles and introducing really applications that we've left to work around them by providing what's called application data gateways and NAT devices that are application-aware to enable applications to work through them and that way itself has really put hurdles in the face of introducing new applications.

So what we're hoping that by adopting IPv6 in certain cases, it really opens the door for enabling new

applications that would work, enabling new services that takes advantage of the access of NAT in the path of communication between any two devices.

Now to get there, so this is really kind of the nirvana or the dream to get to an IPv6 Internet, to IPv6 network, I mean when we talk Internet is really -- I've heard that from Leslie, I heard it from Vint and from others, it's really a network of networks and that's tricky. So there is a network and there is networks that are connected together and those networks provide access to users, access to applications, to application to platforms and all of them have to really work together.

So by definition, it's an ecosystem that's composed of multiple things. It's composed of routing and switching infrastructures, of security gateways that provide firewall services, of load balancers and other things. It's composed of end-systems that operate web services that provides for people who are interested in IMS-based communication, IMS services for establishing VOIP calls, for establishing video conferencing calls as we grow into that in the future.

All that stuff has to work along with each other.
So from personal experience, really, I mean, when we look at IPv6, every time that we used to talk about IPv6, we fill an RFP, we pick one for IPv6. For a lot of people that seemed to be just as a requirement. You could answer it. That's, yes, we plan to. It's not getting into the way of implementing networks or rolling out services because today we're doing based on IPv4, but it verifies whether IPv6 works or not. That's the truth. That's what we lived with for the last few years at least. The reality comes now when you start to actually rolling out the IPv6 or when you start testing IPv6 capabilities on these platforms and you start getting surprises that sometimes what was answered on paper was not really what's happening in the product or what's missing, there is a long-term plan to do it or what thought was to be implemented in hardware is no longer implementable and these are the real challenges that we face today.
What does that mean? That means now we have to really get equipment that provides future parity with what we do today with IPv4 to really be able to roll out IPv6

easily.

How do we get there? I think this is not going to be an easy path. I mean, there will be challenges. I think the equipment we get today are IPv6-capable. That's a fact. Are they complete? No. What the message that we drive to our vendors, at least, today and I hope everybody does, as well, that IPv6 is here. We're not going to get -- especially with new equipment, we're enforcing the message that everything has to be -- either have the IPv6 features already implemented or we have to get a guarantee that the hardware we are acquiring will have IPv6 capabilities that are able to -- whereby the vendor is able to develop IPv6 features that will work on these same hardware, so that we don't have to throw away hardware and acquire new hardware.

The challenge that we still face, that people tell you that I implemented this feature and that feature and this feature and the other feature, what's important really is to implement a complete feature set that enables us to roll out service.

Having features implemented in the individual usually

is not a solution. It's really the set of features that have to come together to enable us to roll out services and that's really again what we tried to drive into the industry.

Now, as the power of the market, how it pulls in IPv6, back in 2007 we had small test bed where we're trying to prototype IPv6 networks and IPv6 services and just because as a company where we come from, we provide voice services.

So one of the things that we're looking at is can we enable VOIP over IPv6. Lo and behold, even the major vendors in the voice space did not have any IPv6 capabilities. The only supported IPv4 hardware did not support it.

Where does the market power come in? A couple of years later, wireless comes in, LTE is on everybody's horizon, IMS is core to a lot of LTE deployments for providing VOIP services, and the decision was to go on IPv6 for these services even, forget about Internet access, even as all garden application.

Guess what happened? Every one of these vendors are going to be a player and still have the market in a

growing economy which is capitalizing on wireless roll-out. Even the same vendors that didn't have it before, they had it and they have it ready. So that's really the demand that pulled it in and really we hope the same thing will happen across all the other equipment from infrastructure, from platform point of view.

Part of the ecosystem is us, where we come from as a network service provider. There are other vendors that we talked about, but there is a content - and I like the analogy that people refer to and we refer to a lot of things in life like that, is the chicken and egg? problem. I provide IPv6 access. IPv6 basic access to what?

All right. You're a content provider. You're providing content to be available to IPv6. The question is who's going to access it? So if we keep asking these questions and these get in the way as hurdles to rolling out IPv6, we'll never have IPv6. We'll have to move into the deployment of IPv6 from the content provider space, from the network service provider space, enable platforms and vendor equipment

routing switching to support IPv6 and synergy, coordinated synergy, with the goal to have an IPv6-enabled Internet because it's really about the network of multiple devices, multiple application platforms, the networks that we're looking to build to really enable efficient and effective communication and that's what really we think we need to drive towards as a community.

>> MR. CHOPRA: Well, thank you. I'm going to do a quick audible, then have you [Vint Cerf] go last, if you don't mind. I want to reserve time for the audience to ask questions.

So what I'd like to do now, if you don't mind, is Peter Dengate Thrush is here from ICANN. Welcome, my man. You're looking very dapper. I like the pink tie. You're on.

>> MR. THRUSH: Thanks, Aneesh, and thank you, Larry, for the invitation and to Jane and her team for all of the work.

Firstly, my apologies for being slightly late. Not intended as any disrespect to the panel or to the topic, simply my naiveté in dealing with D.C. traffic.

I was told it would take an hour to get here which seemed like a lot of time but in fact it wasn't enough. So my apologies.

It's also a pleasure, I have to say, to share the panel with quite a few ICANN people. Obviously Vint's a former chair of the Board of ICANN and George Conrades is an early director of ICANN. George, nice to see you again. And Ram is a liaison. So you really have quite an assembly from ICANN here.

I thought I'd just talk briefly about -- and I'm sitting upstream of John. That's really where ICANN sits in relation to this. It's about the allocation of the IP addresses and just a quick reminder that ICANN runs the IANA function which manages the global pool. This is where they start. This is where they come from for both IPv4 and IPv6 addresses and then there's the five regional registries around the globe, based roughly on continental distinctions. So there's one for Africa, one for North America, Latin America Caribbean, and NCC for Europe.

They are the ones who allocate the addresses in the first place to the ISPs and on a couple of occasions

to actually some countries directly. There are some NICs, some national allocation mechanisms, and they give those out to the customers and I suppose you covered the fact that the /8s that are given to the RIRs have been used in their own -- they have their own policies for allocating those and we're now down to the stage of a policy for what to do when we run out and the policy is that when we get to the last five /8s, they will be given one each to those address registries and then they'll be gone and so that leads to an interesting exercise for ICANN and for the RIRs because there will be a period where there will be a fairly visible market in the IPv4 address space equilibrating at a price just below the cost of installing IPv6. So there's a whole lot of interesting legal and contractual and other issues there.

That allocation of the unallocated ones, the expiring data is, as we see it, some time next year, and all I hear from people is it gets earlier. It started off being at the end of 2011. The last advice I got, though, was it was going to be June or July of 2011.

So it certainly is coming and anyone who thinks this is, somebody said earlier, another rerun of a Y2K exercise we need to get that very clear. This is going to happen.

The good news, of course, is that we've been preparing for it for some time and the RIRs get the IPv6 addresses on exactly the same kind of basis from ICANN, from the IANA function, and allocate them, and as John has explained, they have policies just the same way for allocating those, slightly larger numbers are allocated of IPv6s, of course, than the 4s.

We've been preparing for this for some time. Just some of the historical data. July 1999 was when the first allocations were made of IPv6 addresses to the RIRs. July 2004, glue records were put into the root for .KR, that's the Korean ccTLD, and for .jp, and in October 2006, again the RIRs were given other large blocks of IPv6 address space for allocating, and in February 2008, the first root servers got IPv6 glue [records].

So we've been taking steps to make the network, if you like, really for this and I've got some slides if

people want to see some of the uptake in some of the countries of some of the allocations.

Some other steps that might be interesting. There are some places that we regard as 100 percent IPv6-compliant. You might be surprised to know that they're some of the smaller islands where they now have an equivalent amount of IPv6 as IPv4. One of those is the island of Jersey. Another one which is doing quite well is Cuba which has 75 percent of its space. In other words, it has three allocations of IPv6 and four of IPv4.

Oman has got 50 percent, but it's -- some of them are slightly more serious. For example, my own country, New Zealand, is at 18 percent. Some countries are taking this relatively seriously and some reasonably good drives to get IPv6 allocation implemented. Some of the mechanisms they use -- well, the most significant mechanism that they use is strong association with government, including in procurement programs.

Some of the others, the Czech Republic 19 percent, the Netherlands 17 percent, Malaysia 17 percent, Venezuela

16 percent, Taiwan 15 percent. So other countries are actually making a substantial effort to implement these.

I know Ram talked a little bit about some of the ones that Afilias is associated with. So all together, there are 30 of the top-level domains now signed, more coming on at a regular pace, and we're working at the same time with installing the security system enhancement with watching to see what happens to the root as these changes occur.

I suppose that's probably it, without putting up some graphs. I guess the message from ICANN is that we're ready and the system is ready and now it's up to the community to implement through systems like this the actual mechanisms for industry adoption.

Thank you.

>> MR. CHOPRA: Thank you. We now want to turn to Godfather Vint Cerf to share his thoughts and, Vint, do you mind if I throw a question to you because you're obviously going to have the thoughts on a variety of issues.

>> MR. CERF: Yes, go right ahead.

>> MR. CHOPRA: I'm hearing three deliverables coming out of this session. Deliverable Number 1 is that in 90 days, George - Rock Star that he is - is going to help lead a coalition of the willing to help think through what this board template, risk template might look like, and we'll support this effort to the fullest and I'm confident we're going to deliver some powerful results. So Checklist Number 1, a 90-day development of the risk template. Your reaction to that would be useful.

2. Ram kind of provocatively raised the question of whether or not our engineers are raising the right questions as they deploy. So it's one thing to say we're moving down this road and it's quite another to do the least innovation destruction versus innovation-enabling.

So I don't know if there is a thing which is like the template of the right questions to be asking to make sure that you're going down this path well. I don't know what that is but it seems like it's some kind of a questions template on the engineering side, and I'm hearing a lot of data.

Ram gave some data about the really lame rate at which we've adopted -- maybe lame is the wrong term but --

>> MR. CERF: No. That's a good --

>> MR. CHOPRA: Whatever the number is, the term, and then both James -- I'm sorry -- Peter and John gave statistics and so there's sort of a generic question I have about the data that measures progress and whether or not we're in an open and transparent way seeing this thing happen.

So I see this board deliverable. I see this engineering questions deliverable. I see this transparency kind of dashboard.

Reactions to these things that would help us move the needle and, more broadly, what your vision is to get this thing moving.

>> MR. CERF: So let me try to react first to the three questions you posed.

You saw my hand go up with regard to the checklist. I absolutely believe that's essential and at Google, we have such a thing. So we should compare notes, George.

Second. With regard to this question about what the

engineers are thinking, we've encountered this, too. I mean, we're very far along at Google in the implementation of IPv6.

You had asked three questions of us and I'll try to respond. First of all, I'm here representing Google. We are an applications service provider and I would say that we are probably, on a scale from 1 to 10, we're probably around 8 when it comes to completion of our implementation of IPv6 but this has been going on for almost three years. It has not taken a huge number of people. What it has taken is persistence and meticulous examination of code to find all the places where the assumption is made that v4 -- that an IP address is 32 bits long or to find places where there were burned-in v4 addresses because we didn't need anything else.

So that's taken a great deal of meticulous work and it continues. This is not necessarily the thing you just do once and you're done. Even just as late as yesterday, we encountered a couple places in our v6 implementation that had things that needed to be changed because of assumptions that weren't quite

right.

So it's a continuous process, but it hasn't taken thousands of people. In fact, it started with three engineers who started just working their way step by step through the software and it will continue.

I can't give you a precise date when everything is done, but we have been serving IPv6 for quite awhile now, and I want to come back to that, though, because we've been very selective about who we serve with IPv6.

I'm sorry. The third question?

>> MR. CHOPRA: How do we close the gap?

>> MR. CERF: How do we close the gap? And I think that's what this panel is partly about. So let me try to address that in my planned remarks.

I had -- earlier when Larry was speaking and then before we even started, there was this question of how the hell did we get here and it's my fault. It honestly is.

In 1976, I came to Washington to run the Internet Research Program for the Defense Department and I haven't left. I've been here for 34 years now. And

at that point, we had a couple of implementations, maybe three or four implementations of TCP/IP, but it was like the second iteration of these things.

It wasn't even until November of 1977 that we had three networks actually demonstrate the use of TCP/IP all at the same time, you know, communicating with each other.

So during the course of the 1976 year, the question was how much address space should we allocate for the Internet and there wasn't agreement. One group wanted variable length addresses and they got thrown out very quickly by the programmers who said you'll waste a lot of cycles trying to find fields in the packet if you have to search, you know, to find where they are with these variable length things. So that idea went out the window. Remember back then, computer cycles were expensive.

Then there was another group that wanted 128 bits and to be honest with you, most of the other engineers said you're crazy, why do we need 340 trillion, trillion, trillion addresses? It doesn't pass the red-face test.

Well, then another group said, well, how about 32 bits, that's 4.3 billion addresses. So I'm sitting here with this, you know, bunch of guys screaming and yelling at each other not getting any -- no convergence at all.

So somewhere around 1977, I said, okay, enough, I'm trying to get this thing underway. We haven't -- you know, we have to decide. I said how about 32 bits? It's 4.3 billion addresses. It's enough to do an experiment. We didn't know whether this was going to work. This was an experiment, folks.

The problem is that the experiment never ended and so here we are. It's 2010 and we're now into the production version of the Internet. So that's the basic reason that you're having this problem.

>> MR. CHOPRA: We are having this problem. We.

>> MR. CERF: We are having this problem. Yes, we. Okay. Right.

So let me just -- since others have already touched on a lot of things, I want to do two things, Aneesh. One, I want to say a little bit more about Google and then I want to say something about the more general

problem set that I think still needs to be addressed. No pun intended.

With regard to Google, we absolutely have concluded that we have to run both v4 and v6, so we're a dual-stack operation. We encountered some of the same problems that were described earlier with, for example, the net-scalers and other kinds of load-balancing devices didn't do v6.

We are fortunate in one respect because we make an awful lot of our own equipment. We build our own computers for our data centers and we build our own routers internally in the data network, but we do use commercial equipment, as well. So we have to exactly ask does it handle both v4 and v6 with comparable feature set, which is what Nabil pointed out, I think very properly, that we want the v4 and the v6 networks to be comparable in terms of their functional capability.

So that's required a lot of testing and that leads to this question of when you're acquiring equipment, what is your checklist? What questions are you asking? And so we have been taking time to train our

engineering people to know what it is that they need to be asking in order to make sure that both v4 and v6 are properly accommodated.

We are also now encouraging our engineers to make IPv6 a part of their - what we call - operational key results and these are quarterly targets that each engineer adopts and is then tested on at the end of the quarter to see how much progress has been made. We're making sure that v6 is a part of that vocabulary and that management is seen to be supportive of it because, you know, people are motivated to do things that the managers tell them they'll be rated on to do and if they aren't told that's important, then they don't pay attention to it. So all of that's going on at Google.

I want to come to the question of presenting IPv6 to the users of the Internet. We've been very careful about conditions under which we will actually exhibit an IPv6 address and you all understand that when you go to Google, www.google.com, we actually do something similar to what Akamai does. We do -- we compute a response for you. We don't just hand a table response

back and part of that is that there are cases where what -- you might be in another country and you might want the Google web page to show up in Cyrillic or show up in some other form and so we do that computation.

We are therefore capable of responding with v6 if that seems to be appropriate, especially if the query came in from a v6 source. The question is under what conditions should we in fact respond with a v6 address and the answer is that we'd like to be assured that the party who is getting the v6 address from the DNS lookup is capable of reaching us with IPv6 and this leads to the awkward problem that when the Internet was first being built, everything was connected and the only way it could be connected was with IPv4 and so you grew by accretion and everything was connected to everything.

The IPv6 system is not growing by accretion. It's growing by small bits and pieces. The point that was made earlier about the Internet being made up of a network of networks means that not every network has to implement IPv6 at the same time.

The consequence is that you get islands of Ipv6 connectivity, but you don't have any guarantee that because you're doing IPv6 and some other network is, that there is a path between the two or that the path is adequately performing.

Tunnels, for example, which can provide connectivity are also potentially very fragile and when they break, it's sometimes very hard to figure out what went wrong. So we're not big fans of tunnels. We're fans of having native mode, IPv6 connectivity between our servers and the customers that are trying to reach our applications.

So we've, in our case, implemented a kind of white list that says we won't respond to a v6 query unless we're assured ahead of time by testing that there is going to be adequate connectivity and performance.

This leads me to make a suggestion about what we're going to have to do in order to assure a connected IPv6 network and that is that a liberal kind of interconnection policy among the ISPs. This doesn't have to go on forever but I would urge that serious consideration be given to allowing IPv6 connectivity,

even if you might not have allowed it under v4 rules. There are -- you know, the decision, for example, to do peering or transit and the like is an economic business decision, and liberal interconnection, I think, is going to be needed to assure that we have full connectivity wherever v6 is implemented. If we don't do that, it will be very hard, I think, to have a smooth transition. So that's one thing.

The second observation I would make is that there's -- outside of the service provider world and the application provider world, there is this vast array of users, you in the audience and we on the panel represent a portion of that user space, and our equipment at home may not be IPv6-capable and it's a little hard to see what the motivation will be to run out to pick up the latest, you know, Linksys or what-have-you devices in order to be IPv6-compatible.

Well, I remember that there was a program that was very successfully prosecuted called Cash for Clunkers and now we have some clunky Internet routers and firewalls and things like that. Maybe we should seriously think about cash for clunky routers.

Having said that, I want to remind you that a similar tactic was adopted, Larry, you're familiar with this, when the transition to digital television broadcasts was undertaken. For people who didn't have digitally-capable TVs, equipment was made available and some subsidy was provided.

So I'm -- although I don't pretend to understand what all the implications are of what I just suggested, I'm -- except for the fact that it's going to cost money, I honestly think it's worth at least thinking a little bit about the consumer side of this transition and trying to find ways to encourage it to happen.

Let's see. The -- I have two other points I'd like to make, Aneesh, if I could.

>> MR. CHOPRA: Yes, sir.

>> MR. CERF: One of them has to do with the prevention of, or resistance to, hijacking of Internet address space.

Today it's possible for someone to simply announce that they are connected to a piece of the IP address space and if that information is propagated through the global routing tables, people will route traffic

to that target, even if the target hasn't been assigned that IP address. So that's hijacking. Something called RPKI, which is an attempt to allow the regional Internet registries to digitally sign the assignments of IP addresses, so that people who are doing routing can verify that the party announcing the address is actually authorized to do that is being implemented and my strong recommendation is that it be required for IPv6 allocations so that we have a basis on which to resist this kind of hijacking.

Theoretically, the hijacking should be less attractive in v6 because it's so easy to get address space, but nonetheless it's an issue.

While we're on that topic, there is -- someone mentioned earlier, I forget who, that there is a growing economic -- economics of scarcity showing up with regard to IPv4. People either openly selling IPv4 addresses or a gray market or even a black market evolving for that.

Of course, the problem is if you get an IP address, it doesn't mean anything unless it can be routed. So if it doesn't show up in the routing table, it's a

worthless investment.

What's critical here is that by breaking up the IPv4 address space and trying to sell it in pieces, for it to be useful the routing tables have to get bigger and this is not a good thing because routing tables are already significantly large, and as we move into the IPv6 space, we have to have routing tables with both v6 and v4 entries in them.

So this tendency towards an attempt to either monetize unused IPv4 address space or to somehow buy your way out of having to implement IPv6 by a certain deadline is not a good thing generally for the Internet.

The last thing I'd like to mention goes along with something that you mentioned in your third point, Aneesh, and that had to do with understanding where we are in this whole process and some of the people who are on this panel are in a very good position to provide information publicly about the current state of IPv6 implementation.

In the case of Akamai, because of their quarterly reports that come out about the state of the Internet, which I want to publicly compliment them on, both for

producing and for sharing, you have a potential opportunity to tell more of the IPv6 story. You have a large number of devices that are interacting with each other all the time. You could tell something about how much IPv6 address demand is there. You can tell what kind of connectivity there is at the IPv6 level among the various devices that you use for content distribution. So you're in a very good position, George, to contribute there.

We could probably say something about the v6, the level of v6 queries that we see coming in to the Google environment and I'll go back and ask for that kind of contribution, as well. So building a dashboard to track IPv6, I think, is important, especially from the policy point of view, because it's, I think, important for you in the position you're in, Aneesh, to be able to say something about our national level of IPv6 readiness and use. So I will commit to trying to do everything I can to contribute to that.

So I think that's all I want to take up time with this morning. I'm very eager to hear what the participants

in this workshop have to say, what questions they have and maybe what opinions they'd like to express, and I appreciate the time to address you, as well.

>> MR. CHOPRA: Vint, can I ask you one question before we go to the -- well, we're going to be -- I'm sort of messed up on time on this, but the connectivity policy, the notion that they should be liberal, is that a thing that people could know whether in fact Comcast and Verizon and others have adopted? Is that an artifact that is -- I was teasing over here a little bit.

But how would you turn that into a reality ask? Is there a deliverable; that is, you know, a letter or something signed by the stakeholders that kind of describe what you're asking for?

>> MR. CERF: Well, let's see. I'm going to try to put on a hat that I used to wear because I was at MCI which has now been acquired by Verizon and was responsible for our interconnect and peering policies. To be honest with you, I would not recommend that you force people to make public their criteria. It's a business balancing act to decide whether and how you

peer or by transit or do some other kind of interconnect tactic.

What I would say, though, is that encouraging people to take seriously the need for an IPv6 connectivity and steps that can be taken to achieve it would be a very reasonable thing to advertise, but I don't think I would want to put people up in a microscope. I would allow them to make their own business decisions.

>> MR. CHOPRA: So with that, what I'd like to do is ask a couple of questions from the audience and then do we go three minutes into overtime, Larry, if we're okay, and then we -- yes, sir? And when you ask, if you could please tell us who you are and the organization you're with so it's helpful to the panelists in their response.

>> MR. JACKSON: Bill Jackson with Government Computer News and following up on an issue that was touched upon briefly, how adequate or how mature is the ability of security tools to support IPv6?

>> MR. CHOPRA: Who would like to take that question?

>> MR. CERF: Well, I'll start but I suspect other people will have things to say.

There's no reason why the v6 couldn't be just as secure, if not more secure, than IPv4, but the issue will be whether the things that we use today for security, firewalls, for example, have implemented it. Technically, there's no difference in the functional capabilities that are needed in order to achieve parity with IPv4, for example, in the case of firewalls. So we're back, however, to how much has been implemented and for that, you need to talk to the people who build that equipment.

>> MR. JACKSON: Could I ask another question on that because I think I know the answer, but would part of the problem be the richness of the v6 database versus the v4 database we have today in terms of enabling better security analysis for v6?

>> MR. CERF: I don't -- when you say database, George [Bill], I'm not quite sure what the content of that thing is that you have -- that you're imagining in your head is. So if you're saying knowledge about an IP address from which bad stuff emits, --

>> MR. JACKSON: Yes.

>> MR. CERF: -- for example.

>> MR. JACKSON: Yes.

>> MR. CERF: Well, the answer is that you have to discover the bad stuff when it shows up. We don't have *a priori* reasons to say that particular IPv6 address is a bad address and since there's no binding between v4 and v6 particularly, we don't have a clue based on any v4 address information where v6 attacks might originate.

>> MR. CHOPRA: Thank you, Vint. Leslie had a comment.

>> MS. DAIGLE: Yes. Thanks. I'd like to follow-up on Vint's answer and point out, emphasize the fact that really it's a question of operationalizing IPv6 and, you know, essentially getting -- making sure that your deployments are actually solid operationally which I think is part of what Ram touched on earlier, and I'll take this point to observe that -- to follow up on the point of this not being like Y2K because, you know, whether that's buried in obscure parts of code or not, it is certainly very true that there are network management software tools that aren't quite caught up to supporting IPv6 which is a critical piece

of securing any network, regardless of IPv4 or IPv6.
My point earlier in saying that the v6 climate was not like Y2K in that regard is because we're not yet at the point where we're only worrying about the small bits of code embedded in things and that would actually be a good place to get to.
Thanks.
>> PANELIST: Two quick final comments, if I may.
>> MR. CHOPRA: Sure.
>> PANELIST: One quick comment on the security. I think today we're all very big users of e-mail and one of the primary anti-spam tools that's available for mail server operators is a real-time block list which contains IP addresses to block.
I'm not aware of any of those major vendors that have added IPv6 addressing or have plans to do so in the near term. So that's a potential issue that folks should be aware of and until such time as those tools do exist, I think many mail server operators will be hesitant to accept inbound mail over IPv6.
>> MR. CHOPRA: Comment, Ram.
>> MR. MOHAN: One last thing to add to this is the

fact that, as Vint said, it's not so much about v6 versus v4. Really, the focus from the security perspective has to be about ensuring if you're an organization procuring security solutions, ensuring that the performance capabilities are going to be equivalent because the load that you have to process is potentially far greater and you have to ensure that the performance is as good, if not better.

>> MR. JACKSON: That essentially is my question.

>> MR. MOHAN: So it's not a security question. It's a performance issue.

>> MR. JACKSON: Right, right. And that was the question. Is that performance there? Do they support -- do the security tools support v6 adequately to give you the performance you want?

>> MR. MOHAN: Well, your mileage varies on that.

>> MR. CHOPRA: I think that we'll do is -- yes, one comment.

>> PANELIST: I was just going to say I think it depends really. I mean, it's not something that you could retrofit at some point in time. I mean, the design of the equipment time, the cycles are taken to

exercise given the security rule or whatever it is has to be taken account -- into account to be able to get the performance that was referred to. If it's not taken into account, then you're going to get performance degradation as you apply whatever measure you have for IPv6 and that's really for the equipment design.

>> MR. CHOPRA: I think to be respectful of the next panel, if there's one last brief question, I'll take that. You were first on the hands, so go for it, brother. Scream. We hear you.

>> QUESTION: From MITRE Corporation. We're the FFR in D.C. to help government agencies on the IPV infrastructure and I'm from the Center for Connected Government. So I'm helping multiple agencies. Sitting next to me is my sponsor.

So my question is actually adding on to the gentleman from GCN regarding security. We heard about if you build it, they will come, and we heard people say if we come, they will build it, but I think the key is that -- is on the vendor side.

So from the panel, especially for the people, like

Vint, my question is that have you actually asked vendors that for your next procurement, the IPv6 is the key differentiating factor to decide which vendor you will select?

>> MR. CERF: Well, the answer is certainly, from my point of view, is yes, if we're going to buy anything in the network space, it better be IPv6-capable because otherwise I'm not interested.

>> MR. CHOPRA: Anyone else want to comment on the procurement cycle? Go ahead, please.

>> PANELIST: I'm going to comment briefly but completely.

With respect to firewalls, load balances and security products, we have a case where the industry is looking for those products now about the same time the government is.

So at the same time as the government is now saying how do I become IPv6-enabled, the industry is looking for those products for its live websites.

>> MR. CHOPRA: What a segue to the next panel.

>> PANELIST: With respect to infrastructure, like routers, that actually isn't what happened. In 2005,

the U.S. Government mandated the use of IPv6-compatible equipment which actually led router vendors who would not necessarily have had a market to invest and make this equipment because they had to. They had no choice but to comply with the requirement to make IPv6-IPv4 both in their router equipment.

Now, as the industry has suddenly had to deploy IPv6 infrastructure, the ISPs, the content providers, routers, they've actually been there because the government was requiring them several years earlier in procurement.

So we have a case where the good news is that the underlying infrastructure -- I have not heard people complain about router performance issues. I do have people on the edge of the network saying I can't go there yet because the vendors haven't heard enough demand. This time around, it looks like both industry and government will be asking for that about the same time starting now.

>> MR. CHOPRA: All right. Summary. We're going to have a 90-day turnaround for a template for board risk evaluation. We're going to be on it.

Second. There's some combination of asking the right questions from an engineering standpoint, plus training associated which is a little fuzzy in terms of what it is, but it's in the realm of training and engineering support. So we'll work on that.

And three. Vince kind of threw the ball into our lap, which is how do we establish transparency and, yes, under the policy construct, but the notion is it's privately sourced, publicly available information, not to put too much pressure on George, but some of this could be in his State of The Internet Report.

Whatever the case may be, that's our homework assignment. I hope this was a productive session for all of you. We're moving forward and I want to thank the panelists for their really very productive comments. Thank you for your time.

[Applause.]

>> MR. CHOPRA: All right. It is -- I'll be the last to get up from the group because it is my honor and privilege to introduce to you the moderator for the next session, my brother-in-arms, my partner-in-crime, the Chief Information Officer for the United States -

Vivek Kundra.

Please welcome Vivek to the table.

[Applause.]

>> MR. KUNDRA: All right. Good morning. How's everyone doing?

Well, I'll pick up right where the last panel left off and I think John was talking about the demand side of IPv6. Well, what I'm excited to be here for is one of the key issues that we're trying to address which is execution when it comes to actually deploying IPv6 across the Federal Government.

I am joined by a distinguished set of panelists here. Pete Tseronis, who's the Chairman of the IPv6 Task Force, Doug Montgomery from NIST, and Ron Broersma from DoD, and what we're excited about today is to unveil the guidance that OMB just issued around specific milestones in terms of IPv6 deployment.

Number 1. What we're requiring is that agencies across the Federal Government upgrade publicly-facing, external-facing servers and services by the end of Fiscal Year 2012.

Number 2. We're requiring the upgrade of internal

client applications and communications infrastructure to IPv6 by the end of Fiscal Year 2014.

Number 3. We're designating IPv6 managers across the Federal Government that will be led by the IPv6 Task Force with a focus on execution in terms of migration.

And lastly. Our continued focus when it comes to procurement, to make sure that we're ensuring that agency procurements of network IT comply with the FAR requirements of use with the USGv6 Profile and Test Program for completeness and quality of their IPv6 capabilities.

So we're excited in terms of making sure that the Federal Government is leading the charge when it comes to moving towards operations rather than just strategic plans and as part of the implementation in this guidance, what we're doing is we're going to be conducting actually accountability sessions with a very detailed schedule, so we could unearth some of the challenges that agencies face and bring together a team of technical folks from across the Federal Government that will act as a peer review team in terms of what needs to be done specifically to achieve

these milestones.

Why don't I open up with the first question to Pete in terms of how will the IPv6 Task Force actually help agencies in achieving these objectives?

>> MR. TSERONIS: Okay. Great. Well, it's good to be here and I'm pretty fired up after about four or five years now of watching this progress.

I was just telling Vivek that literally a minute before this, my Blackberry started buzzing because the guidance went by the way of the CIO Council, got to my CIO and much like back in '05, for those of us in the room who got the memo saying please handle this when it was first introducing v6, I got it again. So I just wrote back, well, I'm here with Vivek right now and the Task Force had a role in developing this memo. I'd like to just start out and say that this wasn't done in a vacuum. The champion here in Vivek as well as in Aneesh and the panel we saw before, this is much different than five years ago or 2005 when the memo first came out.

What we're looking at here in terms of trying to accomplish within the Federal Task Force, the IPv6

Task Force, is not only doable but it's something that will involve throughout the Federal Government agency involvement.

Vivek asked about the Federal IPv6 Task Force. I noticed the pat on the shoulder, but, you know, it's not just me. Okay? It's Doug Montgomery from NIST. It's the agency transition managers who we've worked with in the past.

One of the milestones is to have an assignment of transition manager effectively, I think, by October 30th. So there are going to be folks at varying levels in this -- within this task force to take this to the next level and that's June 30, 2008, was a big deal in '05.

We reached it. I think most people will say it was a checkmark. I did pass traffic on a segment of my network that I'm no longer using and I tested it and I got credit for it and green on a scorecard somewhere and that while it seemed watered down, we energized the Federal Government on what's v6 and why we need it.

I don't see us needing to do that again in terms of

raising the awareness and the height and so forth. I see the task force implementing the goals of OMB. We will be seeking out support. Obviously, it's not going to be done by one or two people, but it's going to be something that has to be connected in terms of the message across the Federal Government.

I liked the comment from the last group about this public/private partnership. There may be a Federal IPv6 Task Force but that doesn't mean that the engagement with the private sector community, the service providers that we talked about aren't involved.

John Curran and I have been talking for a couple of months. He's been a mentor and an educator and a friend and at the same time is bringing that global perspective to what the government needs to do.

I can personally attest in the last two and a half years, while things have been a bit silent on IPv6 because folks figured, hey, we reached that milestone, I've spoken with Australia, Germany, Singapore, France, European Union, Japan, and China just on what they're doing, but they've wanted me to explain what's

the USG doing.

Now, a lot of it sometimes is sidestepping and doing a dance and dodging a few bullets because we really haven't come out with anything concrete, but as of today, clearly there's some milestones that I think is going to captivate the global community and that's not being biased, but I really think that folks are going to be excited about the action and the items that are taking place.

So the task force is what it is. It's obviously an arm, but it's -- the agencies have to embrace this. Federal agencies can't do the let's get somebody to manage this for us and what do we need to say that we're done. This is beyond that at this point.

And one of the documents that should help every agency out is this Planning Guide and Roadmap Toward the v6 Adoption in the U.S. Government.

This was developed again with folks in the Federal community as well as in the private sector community in 2009. It's on the CIO Council website. The information, in light of today, will be updated a bit just to show the new milestones and so forth, but this

document was created for the C-SUITE. This wasn't created for the OPS person. It wasn't created for the technical bit-head, the propeller-head.

This is -- this was created and broken out into sections to serve somewhat as a bible, if you will, as how do I get started, why do I need to do this? The question was asked of the last panel: why IPv6, what's the killer application, do I really need to focus on that because the Internet seems to be working fine?

This tells you and recommends in some cases not only where to start with public-facing external servers and the like but how you need to sell this process or how you need to sell this message internally.

The chief procurement officers, the chief acquisition officers, the chief financial officers are equally important to this spreading of the word and embracing of this as much as the CIO is and working in the CIO Office for a number of years, and I am that, I work within the CIO Office at the Department of Energy in my real life, if you will, professional life, my job or someone's job at the Department of Energy's going

to be to bring these people together and say you all play a role in this.

Procurement knows -- needs to know what the FAR language that was released in 2009 means, okay, and working in procurement shops in years past, folks buy goods based on requirements that are handed down to them by way of the CIOs typically, but we need that -- what we need to know now is what are those capabilities that we want in v6 in our agency, at our enterprise.

That's where this work, great work that the work's being done out at NIST is doing, is how do we know what capability we need at the federal level to embrace v6? The Roadmap talks about all these kinds of things to think about and that was the essence of what it was created for.

The enterprise architecture community, the capital planning investment control community, again procurement, the security issues to be thinking about, it's 42 pages, but it's a good read, and it's a read that you can reference, and we will continue to update this as things progress versus when's the next memo

coming out from OMB. It's something we're trying to avoid and again it was not done in a vacuum.

This message, this Task Force has evolved over the past four years and we're pretty excited that we've got the Administration behind pushing v6 to this next level, and I want to echo what I heard in that first panel, which was this isn't going to happen tomorrow, Flag Day, Y2K, unfair comparison.

We said for years, and I think it was John who I first met when I was doing a conference, was this is like the marathon. We may be -- we may have been at Mile 1.2 three or four years ago. We may be at Mile 5.6 now. We've got a ways to go, but once it's baked into technology refresh, once it's baked into strategic planning, Vivek's talked about the Smart Grid, the cloud computing, Aneesh, as well, we now have killer apps that we can associate why we need to do this. Every agency cares about sustainability. Every agency cares about security. Every agency's caring about deploying more and more mobile apps. Well, if you're going to do all that, how can v6 not be in the discussion? If that's not the epicenter of these

deployments and these initiatives, I don't know what is.

If you're going to go to the cloud and you're going to start leveraging the public Internet, I would hope v6 is a concern or a question on someone's mind as to have we baked that in and I said to Vivek before I came in, the worst thing that could happen is right now everybody goes off and builds their v6 network and says I'm done, look how good we are. Ron will talk a bit about the successes and challenges that he's faced.

This isn't about creating in a cloud world a million private clouds out there that don't interconnect. This is the protocol that will connect us. This is an opportunity again to bring v6 into the discussion for any transformative work we're doing and I did happen to have -- and this was not sketched, scripted, but I have the actual innovation document that Aneesh put out because I wanted to see what they were doing in OSTP and they talk about innovation and the Internet. So you team this document when you're making the case for v6 through the Federal Task Force in your own

agency and start bridging that community of interest within your own agency. This is about sharing information and getting agencies to start talking to one another so that federal-wide we aren't -- we are one. We are one v6 interoperable inter-network of communication and I hope that going forward, you know, and supporting Vivek and the mission, as long as they want me to, just to help to be a voice of that, how to do things, how to answer questions, and seek out the people who are doing it.

We're only going to learn from one another and with that, I think would be a good segue to the great work that NIST is doing and Ron. So I'll shut up.

>> MR. KUNDRA: So, Doug, you know, Pete mentioned the FAR and at the end of the day, it comes down to acquisition and the Federal Government actually buying infrastructure that's IPv6-compliant, you know, whether it's the networks program or the federal architects that are engineering these next generation platforms, especially around cloud computing.

Talk to us a little bit about the work that you're doing at NIST specifically around USGv6.

>> MR. MONTGOMERY: Sure. This morning's panel touched on a few recurring themes. One, we ended with the thought that we need to ask for v6 in our products, right? We need to create the demand with the vendors.

An interesting question is so what are you really asking for, right? IPv6 is four characters, right? It sounds concise. We have what we think is a minimalistic technical profile for IPv6 capabilities that cites a 150 RFCs out of the Internet, a 150 technical specifications that define various configurations of IPv6 for real vendor equipment.

So - the question about how do you ask. There was a comment about training acquisition folks. The USGv6 Program was -- we were tasked to fill that void, is to create a technical infrastructure that would allow users to be able to ask for v6 capabilities in a way that is expressed in granularities in ways that users would understand in writing requirements and to provide a technical profile that translates that through to the real technical requirements at the protocol level that vendors have to implement. So

we're there trying to create a vocabulary, if you will, to allow government agencies and others who want to use the profile to select basic levels of capabilities and configurations and translate that through to the technical specifications.

v6 is new to most people, right? So understanding the detailed protocol level requirements is something you would never expect on the consumer side to be that familiar with it.

The second question in that, and it was echoed in this morning's panel, is so what's the completeness and the quality of the products? What if you had a wonderful way of expressing your requirements of what you wanted out of vendors, and then they send you a box that has IPv6 stamped on the outside? We bought some of these boxes, also.

So the second leg of the USGv6 Program is a test program. It's a testing program designed to test the completeness and the quality of implementations. It's a program that has conformance interoperability testing but is a program.

In my long history at NIST, you know, I get a little

antsy about test programs, right? It's easy for them to slide off and become part of the problem space and not part of the solution space.

So we went through a fair amount of effort in the USGv6 Test Program to leverage all that was out there in the industry. So we have a standing agreement with the IPv6 Forum and the IPv6 Ready Logo Program to adopt a suite of test suites for IPv6 that was already vetted by the industry, that there's hundreds of registered products to.

We have taken that test suite, working with the v6 ready community, and we sort of down-select and customize it for the places where the U.S. Government profile had very specific requirements. Some places we specify certain encryption algorithms that they weren't testing. There's some things we need, but there's really -- it's a minimal delta off of what was an industry-driven testing program.

We're doing -- the USGv6 Program is running with commercial labs that are only accredited by us. So there are multiple laboratories that are doing USGv6 testing that vendors can go to and we're really

looking for people to exercise this as sort of a one-stop infrastructure so that folks can begin expressing real requirements to vendors and look for tested results, right?

What we want government agencies to do is to ask and what the FAR actually requires is that you express your requirements towards vendors in using the vocabulary of our profile and ask to see the tested results come back and that's the infrastructure that we're providing that tries to plug both that descriptive capability and that quality control capability.

And to address the last question to the previous panel, one of the gaping holes that we were finding in trying to motivate agencies to both acquire and deploy v6 was exactly network security devices and so one of the things that our profile does, the IETF in this space was somewhat lacking, defines a set of technical requirements for IPv6 firewalls, IDSs and IPSs, and the test program tests against those.

So you can go out and ask a vendor to show you USGv6 test results to prove that their firewall at least has

the minimal set of capabilities, and, mind you, we truly are trying to set low bar capabilities, right? These won't necessarily be best of breed firewalls. There's lot of labs that go on to test other things, but many people, certainly we, probably like you, had bought security devices that claimed to be v6-capable and they didn't even make it to the low bar, right? So we're providing this as a minimal level of what we think is required capability that you can ask that vendors show up and show you their USGv6-tested firewall or IDS.

>> MR. KUNDRA: So, Ron, in light of the new guidance, one of the big challenges is going to be to actually migrate to IPv6 and what's been interesting is you've been able to do that at the DoD level at no additional cost and from an OMB perspective that's very attractive.

So tell me how you did that and what can other agencies learn.

>> MR. BROERSMA: Okay. Thank you. Yeah. A lot of it has to do with being at the right place at the right time in history and so I'll cover a little bit

of that and set some context, as well.

When I first joined the Navy as a civilian working at a research lab, I had the fortunate opportunity of being introduced to this box in the corner that was called an RPANEtIMP and it was IMP Number 3 on the Internet. It's like this cool new thing and it was fun to try for a new software engineer and so – RPANET, yes. Did I say RPANET? Oh, ARPANET. I'm sorry, I'm sorry. I misspoke. Number 3 on the ARPANET, yes.

And so, saw the early days where the protocol was NCP, saw transition to a new protocol called IP, and so this is the second time around for me going from IPv4 to IPv6 and it was originally an eight-bit address. So going to 32 bits seemed like a big thing at the time and certainly looking forward to this next transition.

So for me it was been there/done that, and so when IPv6 became the new protocol, it was definitely something we wanted to play with and in the DoD we started building some test beds and playing with the new protocol, writing some of the early

implementations, and it was in 2003 that the DoD CIO issued a memo saying we want DoD to be IPv6-capable by 2008 and because of some of the test bedding that we were involved with in the research community within DoD, we were selected to be the DoD pilot in the form of the Defense Research Engineering Network which is the ISP for the DoD research community and so DREN was given that task and we started aggressively deploying IPv6 wherever we could and basically seeing what works, what doesn't work, where are the problems, and we learned a lot in that process, and we had a lot of motivated people that wanted to make it a success, but there were some challenges back in that time-frame, certainly, many of which have been resolved now. It's certainly a lot easier than it was back then.

In the old days, we had to worry about operating systems that didn't support it and having to compile in IPv6 code from some place somewhere else in the world where someone had developed it and it was buggy and there was a lot of issues whereas now today it just comes with Windows 7. It's on by default.

So it's a different world. It's much easier and for

most people you should just be able to turn it on and it should just work.

However, we've noticed that people have this notion that it's hard, it's complex, maybe it won't really happen, there's security problems, and a lot of those are just perceptions and so we've had to deal with a lot of those issues as we've deployed it to our networks.

So the way we succeeded, I think, at rolling it out with no cost is we realized that it's not something that happens immediately but we, through smart investments, through tech refresh and leveraging existing staffing that was interested and not really asking for any new money because there wasn't any, we were able to just roll it out over a period of five years.

We had the time to do it and so we were able to achieve success and did not hire a single new person, ask for any extra money to do that, and at least within DREN community, it is now 100 percent dual stack on the wide area and there are large enterprises within the community, like SPAWAR, where it is 100

percent dual stack servers, clients, applications, back-end systems. The public-facing services have been IPv6-enabled for many years and we live it and breathe it every day. There's real people that live and use it constantly and so we know where the bugs are and we're able to achieve that just by generating appropriate interest within the organization.

And I think part of another reason for the success was a corporate culture, where internally we got buy-in from the CIO on down, from the engineers, from the web manager, from the operations staff, from the rest of the IT organization, anybody who touched the network at all or had anything involved with it. Everybody was involved in doing their piece and that was certainly key to the success. We didn't have to have any mandates. We didn't have to pay extra money to the different groups to make it happen.

People realized it was the right thing to do and it got done and part of where we are in the DoD community -- you have to understand there's sort of an operational side of the DoD networking world and a research side.

The operational side run networks that serve our Armed Forces, whereas the research and engineering side services the research, development, test and evaluation, the modeling stimulation community, science and technology, super computers, that whole part of the world operates in the research domain, and so it's in that environment where we really push the envelope, are able to take extra risk and push out IPv6, and when people worried about what is it going to hurt, our motto was don't be afraid to break some glass.

Sometimes you just have to get over that worry and just do it and what we were focused on is what are the things that the operational networks are going to need when they get around to deploying this and turning it on because they're very worried about what are the security impacts, especially when lives might be at stake, and so we were there to make sure the stuff worked, it was available, it was secure, and get the lessons learned on how to deploy it, and so hopefully that addressed the question.

>> MR. KUNDRA: Pete, taking off your IPv6 Task Force

lead hat off and putting on your DOE hat on, with labs all over the country, how -- I mean, how are you approaching this problem, and what is the state of IPv6 implementation within the Department of Energy?

>> MR. TSERONIS: Sure. So without speaking for the scientific community, we're set up as a large agency, a disparate agency. We've got a scientific network, super computer network. We have something that's referred to as ES Net, Energy Sciences Network, as well as the Operational DOENet, DOE, Department of Energy Network.

Much like what Ron said, I would assume, and I think if you were to Google IPv6, you'll see that DOE was one of the standard-bearers of adopting v6 and will probably say, hey, we're running v6 for the reasons we need to run it.

From the standpoint of in the CIO Office and operationally, I think, just based on a few e-mails I got, folks are already scratching their head like who's in charge on this one, and that's what we saw, though, about four or five years ago.

I know at Education, where I was the transition

manager there, just getting people in the room, the first thing I did was call together the leadership and embrace and say, hey, you folks need to champion this. We found really early on if you're in an agency, and I think this would be the case at DOE where you have multiple networks or folks that have multiple data centers, whether it's three servers in a room or our data center out in Germantown, you got to get those people who -- and it's not many. It's probably two, three, or four, coming together and ensuring that they are willing to break a little glass but also say here's what I need to get educated on.

We will have a challenge in the Federal IPv6 Task Force, I believe, just right off the bat getting the right person and that I'm understanding and if they're the right people to rally the folks to say here's why we need to do this.

So large or small agencies, I would expect a smaller agency to be quicker or more rapidly promoting to v6, whether it's to outsource their entire data center through cloud computing and saying, hey, here are our requirements and trusting the whole security aspect

which we're -- is the same kind of question we get, but if you're an agency that wants to own and maintain your data and your infrastructure, first question you should ask is, hey, how old is our stuff? It all has a born-on date and an easy sell to let leadership is going to be we own equipment today through tech refresh could be your ultimate plan.

We own equipment today that has a born-on life and it would no longer be supported. So this is an impetus to move to v6.

I would anticipate at DOE, back to the initial question, is we have to get whoever is designated the transition manager, got to get those right people in the room, leverage the work that's been done with data center consolidation to cloud computing, the work that's already been brought together to say what's the state of our infrastructure, and then team that with what's your strategy going forward.

Forget v6. I would assume mobile computing, cloud computing, sustainability, again to drive the point home. We're already being asked to have plans for that. That's leveraging IT. That's leveraging IP.

So it's not like we have to reinvent the wheel. I would hope that there's not going to be a big learning curve or let me understand why I need to do this again but instead don't break what's already been built and the tools are out there. Okay?

None of the documents that we've talked about today, and people get hung up on documentation and authors, there's a line of sight, okay, that's been going on in the previous Administration up to today, and you could take each one of these and know that there's a linkage and how to use the tools, whether it's defining requirements and capabilities.

I would hope that's what agency transition managers who will be that facilitator to the agency by way of the Task Force will embrace and be passionate about, as you can tell you need to be on this subject matter, or else it's -- you can hear a pin drop because most people will think, you know, there's nothing I need to do. The products I need are already v6-capable inside.

There's work that has to be done to turn on those capabilities and the security issues need to be

addressed, much like they have to be in cloud. So I think it's a -- it's got to be this ongoing discussion.

Again, Mile Marker 5 or 6, what I referred to, is where we are. We're not going to be at the end of this probably before, you know, I retire which is 12 or 13 years down the road hopefully, God willing and my health with some smart meter that's monitoring my blood pressure.

But to that extent, I think it's got to be embraced and you've got to have the right people, you know, leading the effort.

>> MR. KUNDRA: So, Doug and Ron, from a value perspective, so we know that the technical work has been done when it comes to IPv6. The challenge has been actually migrating.

What would you say to a program manager who's managing a billion dollar program whether that's around Medicare/Medicaid or Education or Defense, why should they transition? Why should they take the time now to move forward? What is the value for them in terms of transitioning?

>> MR. TSERONIS: So I think the thing this morning that was well made was the theme of business continuity, right, is for years we chased around the question of what was the killer app for v6? Well, the killer app is that the Internet, as you know it, continues.

That's sustaining the growth that we know about, the potential impacts on innovation. I liked the comments this morning about -- when I used to talk to folks about v6 is I always used to say unlimited addresses and tried to see if there was a spark in their eye which isn't -- I mean, there is the just growing what we do. More customers for the ISPs was discussed this morning, the number of mobile phones.

But really go out to folks with a sense of innovation and tell them unlimited addresses, ask them, you know, do you want to provide a network address to your memory locations on your machines. Think about that power.

So there's a long-term strategic value in moving to v6, right? When we get to the point that known growth will be stopped, certainly innovation will be stopped.

The second thing is in the questions about security. I fully believe with raising the bar on security devices and such, we can get the technology there today to secure networks as we know them.

I think a more interesting question gets to be is what's the risks of non-adoption, right. Many things that we know and trust are based upon the use of v6 addresses, right? So there's lots of trust infrastructures that are linked to addresses. What becomes with those when we get to the point where addresses are no longer global?

There's lots of network management and network abuse systems that rely on being able to identify who's using an address at the end of the Net, right? By the time we get to carrier grade NATs, that trail is going to stop at, you know, the City of Philadelphia and you're one of a 100,000 people behind that NAT box, and then there's the subtle costs of the complexity that we're introducing to keep the current model going, right?

The simple NATs we run in our home was about not paying your ISP, you know, 25 bucks more for the 10

addresses you really needed, something simple, right? When you look at the exhaustion scenarios and it's NAT upon NAT upon NAT in serial networks with the private, the transit, and the FAR Net all in private address space, the rigidity and the, you know, brittleness, the shared state failures that are there, the ability to diagnose that NAT will go down dramatically, right, and so beginning to ask yourselves all these second order questions of the cost of non-adoption, I think, is an important thing for anybody who's managing big programs to think about.

>> MR. KUNDRA: Ron.

>> MR. BROERSMA: Okay. So the question why transition? What's the value?

I think for any organization, as was stated earlier, the real question is if you don't transition, what's the impact? It's really to the point where if you have any -- one example is if you have any public-facing content, there -- at some point where a good part of the Internet will be coming from devices that are IPv6-only and the only way they can get to your old IPv4 content is through some kind of a translator

and if you've looked at the translator scenarios, it's not pretty, and they do not scale. There's going to be a lot of network congestion trying to back-haul that traffic to those translator points.

The tunneling and other transition mechanisms we've seen so far are problematic and are the reason for a lot of the brokenness in the IPv6 Internet today and so we just have to make that go away, but there will be a point where there's a lot of devices, maybe by later this year with some of the new cell phone technology, the 4G or LTE technology, where it's going to be IPv6-only. So they ought to go through translators and so the service to the customer is going to be impacted if you are not delivering your content over IPv6 in the very near term.

Also, just as products mature and IPv6 becomes the mainstream, there will be less support for IPv4 at some point in the product vendors. That's not going to happen near term. It may be five or 10 years but that's certainly something that's on the horizon. So you need to get on the bandwagon sooner than later because some of these transitions to do it right

without huge expense takes five years at least because you have to plan for the tech refresh, you have to plan where you make your investment, maybe choose different products if your vendor does not support v6. So doing it sooner than later is certainly wise at this point, and I think if we look down the road, we can certainly see some benefits from deploying IPv6 once we reach critical mass, once most of the world has migrated because if we can get back to an end-to-end model, like we had in the old days, and eliminate most of the translators, I think there's many things we'll be able to do to create a more robust, reliable, dynamic Internet than we have today and so that is what's mouthwatering for some of us because that can kind of look to that future, but there's just a lot of steps we have to take in the path to trying to get there and those investments are worth making right now.

Some of the benefits we've seen already just with our deployments is things like multicast is just easier with IPv6. It's better thought-out, some of the problematic protocols have been eliminated, and so if

you're doing things like voice and video and doing that over multicast, that can be huge for military applications or for modeling and simulation because it's just much more robust doing it on IPv6 and that for us for awhile was the killer app, was just get rid of IPv4 multicast because it was so painful to keep operating. So v6 is a big win there.

I think security can be a big win because, as we roll out IPv6, it's an opportunity to rethink your addressing plan and to think about how you want to do route aggregation or how you might want to be careful about allocating your address space so that you can build access control lists for certain things you want to protect based on bit boundaries in the address space and there's some huge wins to be gained there when you're not dealing with scarcity of addresses and dealing with very small bit boundaries and so there's some huge opportunities there, opportunities for better security because of sparseness of addresses in large address space, so you can't scan the network and enumerate everything just by doing pings to all the addresses.

So there's improvements for better security down the road potentially and so hopefully we can leverage those some day, once this does become ubiquitous, and so those are some of the examples that one might want to think about in doing it now rather than waiting longer.

>> MR. KUNDRA: Great. Why don't I open up the floor for questions people may have?

>> QUESTION: Thank you. First of all, I really appreciate the comments that Ron and Doug made and I'd like to ask a little bit more about network management because that was a fairly big hole in the v6 space for quite awhile, is having a toolset that not only could do v6 but also do v4 at the same time, you know. A failure generates cascades of alarms and other sorts of things and it all has to get sorted out.

So, Ron, particularly for DREN, can you say a little bit more about the state of completeness of the network management tools that are available and were they commercially available or did you have to build them yourself?

>> MR. BROERSMA: Yeah. Good question. We realized

early on there was a problem there because we thought that since a lot of the network management is done over usually an isolated network, private segment that the world can't get to because it's usually very secure and so we thought normally it's using RFC-1918 space, the private space, and so this would be a great opportunity to try to make something IPv6-only because we didn't have to worry about talking to the outside world.

So we launched a project to take our management network to IPv6-only to see if it could be done. It turned out to be very difficult because every time we turned around, there was some feature that was missing in the vendor products, either in the management stations or the devices that we were managing.

Just turning on something like can you SSH into the box? Well, they don't support SSH yet under v6 and we have to get that implemented and now once you get into the box, do they do SMPv3? Well, no, not quite yet. That will come the next year and then does it export flow records that handle IPv6, and, well, that's another feature that we need.

We spent years feature request after feature request after feature request trying to get those capabilities to where at least in one of the vendors we're now very, very close to being able to turn off IPv4 in our management networks and the last thing, at least for that one vendor, is because things like Link Layer Discovery Protocol, some of those things don't handle IPv6 addresses yet or they weren't in the implementation and so we're down to the last piece where we'll finally be able to do that.

But we've noticed some things which are very interesting. Because we carefully chose our IPv6 addresses, we use the ULA-type of addresses, we're -- our addressers are shorter than our v4 addresses, less typing, because of how we allocate the bits.

The other thing which was interesting is because of the routing tables being smaller, routing convergence for the management networks are extremely fast because of the aggregation of address space if it's done correctly.

So if you use careful allocation of addresses in a good address plan, you can do some interesting things

and actually make your network much more robust, but there's a lot more to do. Not all the vendors support all the pieces yet, but it's -- we've pushed very hard on the vendors to provide those and point out the failings and they've been slow but have been coming along and -- but that's all part of the whole feature parity argument, is we've noticed that there's just so many pieces of vendor implementations, that the v6 features that are there are the ones that are documented in the RFCs and the ones that will get them through the profiles and checklists so that they can sell to the government.

But every vendor has something in their product which we consider the secret sauce which makes their product interesting and why we buy their product. Those features don't necessarily have IPv6 support and so feature parity is one of our biggest challenges with all the vendors and the other problem with vendors is we found many cases, we found bugs where it's like “didn't anybody test this?”, and it's because their QA suites have not matured to the degree in their IPv6 spaces. They have an IPv4 and many vendors aren't

using their own products, using IPv6 in their own networks, and so we've identified that as a problem of vendors not eating their own dog food and have really pushed that and publicized competition among vendors to see who's doing it, who's not doing it, and are starting to get some good traction there, as well. So it's a lot of working with industry, pushing the vendors, identifying the missing pieces and getting those resolved and so I think we've made huge progress in the last few years.

>> MR. KUNDRA: In the back over there.

>> MR. CACUS: Thank you. Max Cacas from Federal News Radio.

I'm trying to get an idea here of where we are and I've talked to Pete Tseronis about this a lot in the past. We've been covering this story a lot.

Pete, can you give us an idea right now of where the Federal Government stands on a percentage basis of agencies in terms of IPv6 implementation, who's there, who's in the same ballpark with Ron over at DREN, who still needs help, and how many agencies, percentage of agencies in the government are on their way? What do

we know so far?

>> MR. TSERONIS: Well, we don't have an exact metric, but I can tell you probably by the end of January 2011 we will, based on the latest task force schedule, because we'll visit with all these primary agencies to see kind of where we are.

I can speak from experience. I know for a fact that Education had a very mature plan and after I left, as well. Obviously the case of DREN, you know, engineering network really pushing the envelope. Yeah. What is -- does that encompass DoD? Probably not. Where does DISA stand on their operational deployment? We don't know that. We know that there are pockets out there that are out-sourcing with v6 in mind and really haven't given it a second thought. They've said, hey, I have outsourced totally. So how do you qualify or quantify that agency? Are they v6, you know, in the next stratosphere?

I hope -- short answer is that after we kick this off and again post June 30th, we haven't done much of this analysis because agencies have really been saying we're going to use our enterprise architecture and

tech refresh planning which was the direction from your predecessor to migrate to this next generation Internet. We haven't tracked that. We don't know from a task force perspective if we're even grading people where they are, but the assumption has been that networks are maturing and again I go back to Ron's point of this is that opportunity where people have to go back to the blackboard or the chalkboard and say I have to readdress my network. Ugh!

You know, I mean that's daunting. I've been working in ops for a number of years in my early life. That's not something you want to do. You figure out -- I don't want to call them band-aids, but things like NAT, classless inter-domain routing, and, you know, classful networks, those were opportunities to kind of say, well, let's do that instead of kind of going back to the drawing board.

I think, Max, after these initial discussions, we'll get a sense of where agencies are and we'll be able to report that out because obviously Vivek here wants to get a sense so that we can be very execution-oriented -- that's probably not a good word -- moving forward

as to where agency progress is and to really lean on one another.

A guy like Ron and the success that he's had in DREN at the Navy should permeate throughout DoD. If Education again is advanced, have they deployed and started using applications with their IP Telephony Network which is nationwide? How are they dealing with the security implications of IP?

So to say to the agencies and assume that they're using tech refresh, do they have IPv6 test labs even? I've heard from some agencies saying we can't even get funding to put a test lab together to do this. So how do we buy product and then deploy it if we don't have the money to really test it, much like Ron's been able to do?

>> MR. BROERSMA: To be fair, the question was probably too general because in 2008, I mean, there was a milestone there about getting your ISP services and your core, if you will, transit and gateway to it v6-enabled and I believe all agencies met that goal. But from that point on, the milestones have been very broad and very general, adopt v6. That's very

general.

The new guidance actually has very specific milestones, right? External-facing servers by the end of 2012. That's a much more discrete goal and milestone, right? So the answer to your question might be depends on what you wanted to measure. I think with the new -- with new guidance here, it will be very clear when we achieve those goals and exactly which systems we're talking about upgrading when.

>> MR. KUNDRA: Is there a question on this side of the room? No. Okay. Any other questions?

>> MR. McDERMOTT: My name is Greg McDermott. I'm with a small company called New Dynamics. I don't know if you're familiar with what we do.

We've been on the front-end with just about every one of the network manufacturers to help them certify and accredit their technologies before they go GA.

My challenge here is I'm responsible for building the Federal practice here and the daunting task for what we've seen is that you've got multiple heterogeneous environments where multiple different technologies that have passed the IPv6 muster but when you link

them all together in an overall, you know, implementation, there's those challenges.

Does this new guidance and some of the direction that you do with your programs and what you do encompass that whole end-to-end view and policy?

>> MR. TSERONIS: Is that directed to me? Okay.

Well, I'm looking at the memo right here, okay, and really if you were to look at the previous -- well, I shouldn't say the previous memo, but post June 30th, it was really to going forward the plan was -- and don't quote me on this one, but it was to ensure that you have an end-to-end secure IPv6 network, leveraging technology refresh in your strategic planning, etcetera, etcetera.

That's great and that's assuming folks embraced v6 and said we have a roadmap in place and we're trying to go to this innovative technologically-advanced network which is what IT's all about, right? Evolving.

This one is calling out, you know, specific things that are actionable, that I think, you know, with Vivek's leadership and where he's pushing the envelope a bit, 2012 seems like a far way away. It's not. You

have to start budgeting two years in advance for what you're planning to do if you're a federal agency, right?

2014. I mean, this is forcing agencies, I believe, to think about that end-to-end picture. This is forcing agencies to say what are we trying to do if we even have an IT strategic roadmap. I know we have one at the Department of Energy, having just completed a cloud sourcing exercise, and we know that we -- we know the services that the Department of Energy wants. We now can map services to capabilities, such that v6 would be the protocol that's going to get us there. So I would charge agencies when they read this to say this isn't an exercise of assign somebody and say that you're doing such and such. It's going to cause you -- it should force an agency to say what is our plan for the next five years? What are the services we want to deploy and how is v6 going to get us there? That's the business case and the story for any future investment dollar requests or simply to leverage tech refresh to get where you want to get to, and I got to keep underscoring the point that, depending upon the

size of the agency, that could be an excuse for folks not communicating.

But you're seeing what DREN, as advanced as they are, you hope that the op side of the house isn't saying, well, that's the engineering network. They don't have to worry about the same things we do. You can't have that discussion. You've got to have people in ops with that same sort of, okay, what is it we need to do even if it means going back to the drawing board? My two cents.

>> MR. BROERSMA: Maybe related to that question is some agencies approach the Task Force about potentially leveraging systems integration testing. Our testing program is a product testing program, right? So it's products and IPv6 implementations in isolation, but maybe your question was about how do you get to sort of full systems integration testing and the fact that every agency will have to deal with that and the potential to maybe share that activity and leverage that activity is something that's been brought up to the task force before.

>> MR. KUNDRA: Last question.

>> QUESTION: I have two questions. First is, is this guidance document generally available, and if it is, where is it?

>> MR. KUNDRA: It's going to be on cio.gov.

>> QUESTION: And when should we expect to see that?

>> MR. KUNDRA: It should be up now, if it's not already up.

>> QUESTION: Okay. Good enough. And then the next one is how will the Fed act as an enterprise on this effort?

This is something that we have 18-26 different agencies doing. Why would we all do it 26 different ways? What types of resources or mechanisms, like the task force, will allow us to act corporately?

>> MR. BROERSMA: Well, I would echo -- first off, I bet you Danny has it in his in box right now, CIO over at Education. So you might want to have a talk with him when you get back there, Steven.

Vivek's point initially was what's the Task Force -- we see the Task Force, which I got to be honest, right now it's about four people, right, leadership-wise. We are looking to expand that to the point where we

have got the right people, the right people and then folks beyond the right people embracing this.

So I could see in your case in the Office of Enterprise Architecture that we got to have someone who has that business view of IT as well as an operational aspect of IT as part of a Task Force extension.

First time around, we populated it with vendors and with the folks who maybe were one or two levels down in the data center. It was really hard to sell that to the business focus. So we created that line of sight discussion and while there were just another myriad of questions about why do I need to know this, that's what we're looking this time around is to maybe cull together that mishmash of people who have those two perspectives. Maybe they've worked in two different parts of that organization.

So the Task Force is -- yes, it's back in business but it's undergoing a bit of a shake-up in a positive way to kind of blend, right, the folks, the right folks that need to be -- much like we saw in the panel earlier today. Everybody had a specific role but it

was a common thread there in terms of experience. So we see that being a key aspect.

>> MR. KUNDRA: And a key part in terms of what we've done historically around issuing policy versus a huge focus on execution, a good model here is to look at what we've been able to do with the data center consolidation effort where we essentially brought together a team of 25 people across the Federal Government, very, very focused, very surgical in our approach in terms of looking at the unit of analysis as a department, but bringing external resources to provide the peer review process and then tying that to the 2012 budget decisions, so that it wasn't -- the effort wasn't in the abstract in terms of just strategy or a planning document but had serious consequences in terms of what comes out of that work. So that's part of what we're trying to do here with the migration, is set very specific milestones with actual outcomes, not just in the abstract that everybody should move to IPv6, it's almost like saying everybody should work out because it's healthy, but we're trying to be very specific and then following it

with, you know, best practices.

How do we learn from what is working, what is not working, what are the barriers, and then the key here in terms of the migration, what's going to determine success or failure is going to be a relentless follow-up, not just fire and forget in terms of the guidance but being relentless on a monthly, quarterly basis, looking at this across the Federal Government, and holding the IPv6 transition managers accountable for the performance of their agencies and being very transparent around how we're actually performing across the Federal Government.

One more question and then I think we have to wrap up.

>> QUESTION: First of all, I wanted to make an observation that it was the operations group at Google that initiated the effort to put IPv6 into work as opposed to engineering or anything else. I thought that was kind of cool.

I have one question and one other observation. The question has to do with interoperability testing and it's for Doug. Is there something going on in addition to unit testing that NIST is undertaking or

is there anyone else like University of New Hampshire or some of the others who have done this in the past, Connect-a-thon-like things?

>> MR. MONTGOMERY: It's actually part of our test program [and] is two-pronged. It's conformance inter-op testing, no product that comes out of a USG-accredited test lab, of which the University of new Hampshire is one of the accredited test labs, will not have gone through inter-op, and I should say in the development of test programs, if you think about the fact that I noted our profile sites a 150 RFCs, if you tried to put quality consensus test suites in place for all of them, we wouldn't be done until after Pete retires.

So what we've done is strategically chose -- (1) we adopted the v6-ready logo test suite for as much as we could which gave us a giant bootstrap, maybe covering 80 percent of the profile. At that point we've been on an expedition path of filling in the gaps but any time there's a protocol that isn't covered by a consensus test suite, we write the inter-op test suite first and then the conformance second.

So the real proof–in-the-pudding is inter-op, all the devices that will come out will be inter-op tested. There's other aspects about the test program that we've sort of adopted a full disclosure mode, right, and one of the things that has to be declared in the output of a test suite is, is your device capable of operating in a v6-only mode because we see the same tensions with control playing management protocols and the only thing -- because the government is in the position of now asking or requiring that you buy this stuff.

We don't -- we want to protect the government investment. We don't want you to be blindsided by it. So a product that comes out of one of our test labs will have a report that, if nothing else, you will see that a check box that this box either is or isn't capable of fully operating in a v6-only environment and if no, please explain, and at least you'll know. You won't be blindsided by it, but inter-op is a key leg and maybe the more important leg of the test program.

>> QUESTION: Doug, last comment to that. Your

mention about follow-up, I think, is absolutely on target. But I also like the comment about consequences, particularly budget consequences, as in if you don't implement v6 by this amount of time, you can kiss off, you know, your President's line item in the President's budget or something along those lines.

>> MR. KUNDRA: On that note, well, thank you very much.

[Laughter.]

>> MR. KUNDRA: I want to thank Pete, Doug, and Ron for taking the time to be on this panel and for their thoughtful input and also for helping us lead this transition. Thank you.

[Applause.]

>> MR. KUNDRA: And at this point, what I'd like to do is actually introduce -- if you can give me one second here.

At this moment what I'd like to do is introduce Anna Gomez, who is the Deputy NTIA Administrator, to deliver some closing remarks.

Thank you.

>> MS. GOMEZ: I know I'm what stands between you and

lunch and you and all the rock stars that we've had at today's workshop, but I want to thank our industry panelists, our Internet technical community panelists, our government colleagues from the IPv6 Task Force, and our two moderators, the nation's CTO Aneesh Chopra and the nation's CIO Vivek Kundra. We can't thank you enough for facilitating this excellent event today.

I also want to thank Jane Coffin and Yvonne Neal-Barfield for their hard work in pulling this together along with our other NTIA staff people.

And, finally, I want to thank Fiona Alexander for all of her worldwide leadership on many of these issues.

This event underscores NTIA's commitment to convene stakeholders to highlight important Internet technology issues, particularly dynamic issues like this one, that we heard about today that really impact technology, development, and innovation. We have convened this workshop as the Administration is committed to government and industry deployment and adoption of IPv6 and is facilitating these discussions to highlight this issue.

Today, we've heard that IPv6 uptake and adoption is an

important and urgent issue that can be successfully handled with good planning. In addition, IPv6 is important to innovation, advanced network build-out, smart grid, cloud computing, smart phones, and future applications.

So deploying IPv6 is a critical investment for the continued growth of the Internet economy. Therefore, we encourage companies to continue to share best practices on IPv6 uptake for all businesses to benefit, in particular some of our small- and medium-size colleagues, as well.

We look forward to seeing more of the U.S. Federal Task Force on IPv6 and we look forward to hearing from industry and from the Internet technical community on what progress we are making and, of course, on what progress we are making on our new initiatives that we learned about today from Aneesh's panel.

So please stay in contact with us, with your government representatives here, so that we can continue to help foster this innovation and now you can go to lunch.

Thank you.

[Applause.]